High Dimensional Multi-Level Covariance Estimation and Kriging

Julio E. Castrillón-Candás

arXiv:1701.00285v1 [stat.CO] 1 Jan 2017

January 3, 2017

Abstract

With the advent of big data sets much of the computational science and engineering communities have been moving toward data-driven approaches to regression and classification. However, they present a significant challenge due to the increasing size, complexity and dimensionality of the problems. In this paper a multi-level kriging method that scales well with dimensions is developed. A multi-level basis is constructed that is adapted to a random projection tree (or kD-tree) partitioning of the observations and a sparse grid approximation. This approach identifies the high dimensional underlying phenomena from the noise in an accurate and numerically stable manner. Furthermore, numerically unstable covariance matrices are transformed into well conditioned multi-level matrices without compromising accuracy. A-posteriori error estimates are derived, such as the sub-exponential decay of the coefficients of the multi-level covariance matrix. The multi-level method is tested on numerically unstable problems of up to 50 dimensions. Accurate solutions with feasible computational cost are obtained.

KEY WORDS: Hierarchical Basis, Sparse Grids, Machine Learning, High Performance Computing, Sparsification of Covariance Matrices, Random Projection Trees, Fast Multipole Method

1 Introduction

Big data sets arise from many fields, including, but not limited to commerce, astrophysical sky-surveys, enviromental data, and tsunami warning systems. With the advent of big data sets much of the computational science and engineering communities have been moving toward data-driven approaches to regression and classification. These approaches are effective since the underlying data is incorporated. However, they present a numerical challenge due to increasing size, complexity and dimensionality.

Due to the high dimensionality of the underlying data many modern machine learning methods, such as classification and regression algorithms, seek a balance between accuracy and computational complexity. How efficient this balance is depends on the approach. Linear methods are fast, but only work well when there is linear separation of the data.

For non-linear description of the data, kernel approaches have been effective under certain circumstances. These methods rely on Tikhonov regularization of the data to obtain a functional representation, where it is assumed that the noise model of the phenomena is known. However, this assumption is not necessarily satisfied in practice and can lead to significant errors as the algorithm cannot distinguish between noise and the underlying phenomena.

Sparse grid polynomial approximation [3, 4, 6, 7, 10, 26–28, 31] has emerged as one of the most appealing methods for approximating high-dimensional functions. It is simple to use and converges significantly faster if the function at hand presents some degree of regularity. In [30] the authors develop a sparse grid approach to high dimension classification. However, the noise is assumed to be known, which can still lead to significant error.

To incorporate the variability of the noise model a class of machine learning algorithms based on Bayes method have been developed. In this approach the noise model is assumed to be known up to a class of probability distributions and an optimal choice is made that fits the training data and noise. For example, from the Geo-statistics community a well known approach to identifying the underlying data and noise model is known as kriging [25]. The noise model parameters are estimated from the Maximum Likelihood Estimation (MLE) of the likelihood function.

Kriging methods are effective in separating the underlying phenomena from the noise model. In practice the covariance matrices tend to be ill-conditioned with increasing number of observations making kriging methods numerically fragile. Moreover, applications are limited to 2 or 3 dimensions as the computational cost increases significantly with dimension.

In [8] a novel algorithm to solve kriging problems is proposed. The method is fast and robust. In particular, it can solve kriging problems that where not tractable with previous methods. However, this approach is limited to 2 or 3 dimensions and the computational cost scales very fast with spatial dimension, thus making it impractical for high dimensional problems. Other kriging methods have been developed using skeletonization factorizations [24], low-rank [29] and Hierarchical Matrices [19, 21] approaches.

In this paper we extend the kriging approach in [8] to high dimensions. This novel approach effectively identifies the high dimensional underlying phenomena from the noise in a numerically stable manner. Our approach transforms high dimensional ill-conditioned covariance matrices to numerically stable multi-level covariance matrices without compromising accuracy. In Section 2 the problem formulation is introduced. In Section 3 Multivariate polynomial approximations and sparse grids are discussed. In section 4 it is shown how to construct a multi-level basis with a random projection that scales well with dimension. In section 5 the construction of the multi-level covariance matrix is discussed. In section 6 the multilevel estimator and predictor are formulated. In section 7 error estimates for the decay of the multi-level covariance matrix and the inverse solution are derived. These estimates are based on the Smolyak sparse grid formulas.

Note that for the less mathematically inclined reader this section can be skipped. In section 8 the multi-level kriging method is tested on numerically unstable high dimensional problems. The results are obtained with good accuracy and feasable computational cost.

2 Problem setup

Consider the following model for a Gaussian random field Z:

$$Z(\mathbf{x}) = \mathbf{k}(\mathbf{x})^{\mathrm{T}}\boldsymbol{\beta} + \varepsilon(\mathbf{x}), \qquad \mathbf{x} \in \mathbb{R}^d, \tag{1}$$

where d is the number of spatial dimensions, $\mathbf{k} : \mathbb{R}^d \to \mathbb{R}^p$ is a functional vector of the spatial location $\mathbf{x}$, $\boldsymbol{\beta} \in \mathbb{R}^p$ is an unknown vector of coefficients, and ε is a stationary mean zero Gaussian random field with parametric covariance function $C(\mathbf{x}, \mathbf{x}'; \boldsymbol{\theta}) = \mathrm{cov}\{\varepsilon(\mathbf{x}), \varepsilon(\mathbf{x}')\}$ with an unknown vector of positive parameters $\boldsymbol{\theta} \in \mathbb{R}^d$.

Suppose that we obtain N observations and stack them in the data vector $\mathbf{Z} = (Z(\mathbf{x}_1), \ldots, Z(\mathbf{x}_N))^{\mathrm{T}}$ from locations $\mathbb{S} := \{\mathbf{x}_1, \ldots, \mathbf{x}_N\}$, where $\mathbf{x}_1 \neq \mathbf{x}_2 \neq \mathbf{x}_3 \neq \cdots \neq \mathbf{x}_{N-1} \neq \mathbf{x}_N$. Let $\mathbf{C}(\boldsymbol{\theta}) = \mathrm{cov}(\mathbf{Z}, \mathbf{Z}^{\mathrm{T}}) \in \mathbb{R}^{N\times N}$ be the covariance matrix of $\mathbf{Z}$ and assume it is positive definite for all $\boldsymbol{\theta} \in \mathbb{R}^w$. Define $\mathbf{M} = \big(\mathbf{k}(\mathbf{x}_1) \ldots \mathbf{k}(\mathbf{x}_N)\big)^{\mathrm{T}} \in \mathbb{R}^{n\times p}$ and assume it is of full rank, p. The model (1) leads to the vectorial formulation

$$\mathbf{Z} = \mathbf{M}\boldsymbol{\beta} + \boldsymbol{\varepsilon}, \tag{2}$$

where $\boldsymbol{\varepsilon}$ is a Gaussian random vector, $\boldsymbol{\varepsilon} \sim \mathcal{N}(\mathbf{0}, \mathbf{C}(\boldsymbol{\theta}))$. The aim now is to:

- *Estimate* the unknown vectors $\boldsymbol{\beta}$ and $\boldsymbol{\theta}$;
- *Predict* $Z(\mathbf{x}_0)$, where $\mathbf{x}_0$ is a new spatial location. These two tasks are particularly computationally challenging when the sample size N and number of dimensions d are large.

The unknown vectors $\boldsymbol{\beta}$ and $\boldsymbol{\theta}$ are estimated with the log-likelihood function

$$\ell(\boldsymbol{\beta}, \boldsymbol{\theta}) = -\frac{n}{2}\log(2\pi) - \frac{1}{2}\log\det\{\mathbf{C}(\boldsymbol{\theta})\} - \frac{1}{2}(\mathbf{Z} - \mathbf{M}\boldsymbol{\beta})^{\mathrm{T}}\mathbf{C}(\boldsymbol{\theta})^{-1}(\mathbf{Z} - \mathbf{M}\boldsymbol{\beta}), \tag{3}$$

which can be profiled by generalized least squares with

$$\hat{\boldsymbol{\beta}}(\boldsymbol{\theta}) = \{\mathbf{M}^{\mathrm{T}}\mathbf{C}(\boldsymbol{\theta})^{-1}\mathbf{M}\}^{-1}\mathbf{M}^{\mathrm{T}}\mathbf{C}(\boldsymbol{\theta})^{-1}\mathbf{Z}. \tag{4}$$

A consequence of profiling is that the maximum likelihood estimator (MLE) of $\boldsymbol{\theta}$ then tends to be biased.

To address the prediction part, consider the best unbiased predictor $\hat{Z}(\mathbf{x}_0) = \lambda_0 + \boldsymbol{\lambda}^{\mathrm{T}}\mathbf{Z}$ where $\boldsymbol{\lambda} = (\lambda_1, \ldots, \lambda_n)^{\mathrm{T}}$. The unbiased constraint implies $\lambda_0 = 0$ and $\mathbf{M}^{\mathrm{T}}\boldsymbol{\lambda} = \mathbf{k}(\mathbf{x}_0)$. The minimization of the mean squared prediction error $\mathrm{E}[\{Z(\mathbf{x}_0) - \boldsymbol{\lambda}^{\mathrm{T}}\mathbf{Z}\}^2]$ under the constraint $\mathbf{M}^{\mathrm{T}}\boldsymbol{\lambda} = \mathbf{k}(\mathbf{x}_0)$ yields

$$\hat{Z}(\mathbf{x}_0) = \mathbf{k}(\mathbf{x}_0)^{\mathrm{T}}\hat{\boldsymbol{\beta}} + \mathbf{c}(\boldsymbol{\theta})^{\mathrm{T}}\mathbf{C}(\boldsymbol{\theta})^{-1}(\mathbf{Z} - \mathbf{M}\hat{\boldsymbol{\beta}}), \tag{5}$$

where $\mathbf{c}(\boldsymbol{\theta}) = \mathrm{cov}\{\mathbf{Z}, Z(\mathbf{x}_0)\} \in \mathbb{R}^n$ and $\hat{\boldsymbol{\beta}}$ is defined in (4).

Now, let $\alpha := (\alpha_1, \ldots, \alpha_d) \in \mathbb{Z}^d$, $|\alpha| := \alpha_1 + \cdots + \alpha_d$, $\mathbf{x} := [x_1, \ldots, x_d]$. For any $h \in \mathbb{N}_+$ (where $\mathbb{N}_+ := \mathbb{N} \cup \{0\}$) let $\mathcal{Q}_h^d$ be the set of monomials $\{x_1^{\alpha_1} \ldots x_d^{\alpha_d} \mid |\alpha| \leq h\}$. The typical choice

for the matrix $\mathbf{M}$ is to build it from the monomials of $\mathcal{Q}_h^d$ with cardinality $\binom{d+h}{h}$. However, as the number of dimensions d increases the size of the parameter vector β becomes difficult to handle.

To mitigate this problem we need to control the number of polynomial terms along each dimension while still retaining good accuracy by selecting a judicious set of polynomial functions $\mathbf{k}(\mathbf{x})$.

The second challenge is that the covariance matrix $\mathbf{C}(\boldsymbol{\theta})$ in many practical cases is ill-conditioned, leading to slow and inaccurate estimates of $\boldsymbol{\theta}$. In this paper, we propose a new transformation of the data vector $\mathbf{Z}$ leading to a decoupled multi-level description of the model (1) without any loss of structure. This multi-level representation leads to significant computational benefits when computing the kriging predictor $\hat{Z}(\mathbf{x}_0)$ in (5) for large sample size N and high dimensions d.

3 Multivariate polynomial approximation

We now have to choose an appropriate model $\mathbf{k} : \mathbb{R}^d \to \mathbb{R}^p$ of random field Z. The choice of $\mathbf{k}$ (and thus $\mathcal{P}^p(\mathbb{S})$) will determine the accuracy of the representation of the underlying deterministic function. Indeed, there are two main factors that we should consider. First, the size of p and second, the accuracy of the approximation in a suitable functional space. The balance between these two factors should be carefully considered.

Without loss of generality for $n = 1, \ldots, d$ let $\Gamma_n := [-1, 1]$ and $\Gamma^d := \prod_{n=1}^d \Gamma_n$. Furthermore:

i) Let $\mathbf{p} := [p_1, \ldots, p_d] \in \mathbb{N}_0^d$, $w \in \mathbb{R}$ and $\Lambda(w) \subset \mathbb{N}_0^d$ be an index set that will determine the indices of polynomial basis functions, i.e. p.

ii) Restrict the number of polynomials along each dimension by using the set of monomials contained in

$$\mathcal{Q}_{\Lambda(w)} := \left\{ \prod_{n=1}^{d} y_n^{p_n}, \text{ with } \mathbf{p} \in \Lambda(w) \right\}$$

i.e. $\mathbf{k}(\mathbf{x})$ is built from the monomials in $\mathcal{Q}_{\Lambda(w)}$.

The choice of index set $\Lambda(w)$ will determine the accuracy of the polynomial representation and the size of p. We first consider the index set for a full Tensor Product (TP) grid $\Lambda(w) \equiv \{\mathbf{p} \in \mathbb{N}_0^d : \max_{i=1}^d p_i \leq w\}$. This is a poor choice as p increases $\prod_{i=1}^d (p_i+1)$ and quickly becomes intractable. Another choice is the Total Degree (TD) index set, which grows as $\binom{d+w}{w}$, but still suffers from the *curse of dimensionality*, thus impractical for high dimensions.

Smolyak (SM) and Hyperbolic Cross (HC) index sets are a popular choice for large dimensions. SM have been used extensively in sparse grid approximations of high dimensional analytic functions [27, 28]. HC have been used in the context of Fourier approximations of high dimensional functions. In Table 1 the different choices for $\Lambda(w)$ are summarized.

In Figure 1 a comparison between TP,TD,SM and HC is shown. As observed, the SM and HC grows much slower with respect to the higher polynomial order.

Approx. space	Index Set: $\Lambda(w)$
Tensor Prod. (TP)	$\Lambda(w) \equiv \{\mathbf{p} \in \mathbb{N}_+^d : \ \max_{i=1}^d p_i \leq w\}$
Total Degree (TD)	$\Lambda(w) \equiv \{\mathbf{p} \in \mathbb{N}_+^d : \ \sum_{i=1}^d p_i \leq w\}$
Smolyak (SM)	$\Lambda(w) \equiv \{\mathbf{p} \in \mathbb{N}_+^d : \ \sum_{i=1}^d f(p_i) \leq w\}$ $f(p) = \begin{cases} 0, \ p = 0 \\ 1, \ p = 1 \\ \lceil \log_2(p) \rceil, \ p \geq 2 \end{cases}$
Hyperbolic Cross (HC)	$\Lambda(w) \equiv \{\mathbf{p} \in \mathbb{N}_+^d : \ \prod_{i=1}^d (p_i + 1) \leq w\}$

Table 1: Index set of different polynomial approximation choices.

3.1 Sparse grids

Consider the problem of approximating a function $u : \Gamma \to \mathbb{R}$ on the domain Γ^d. A good choice for approximating such functions are Smolyak sparse grids with an appropriate set of abscissas. Note that in this paper a sparse grid representation is not explicitly used, unless the observation nodes correspond exactly to the abscissa grid. Rather, it is used to estimate the accuracy of a sparse multi-level covariance matrix.

For $q \in \mathbb{N}$ define the following spaces

$$L^q(\Gamma^d) := \{v(\mathbf{y}) \,|\, \int_{\Gamma^d} v(\mathbf{y})^q d\mathbf{y} < \infty\} \ \text{ and } \ L^\infty(\Gamma^d) := \{v(\mathbf{y}) \,|\, \sup_{\mathbf{y} \in \Gamma^d} |v(\mathbf{y})| < \infty\}.$$

Now, let $\mathcal{P}_\mathbf{p}(\Gamma^d) \subset L^2(\Gamma^d)$ be the span of tensor product polynomials of degree at most $\mathbf{p} = (p_1, \ldots, p_N)$; i.e., $\mathcal{P}_\mathbf{p}(\Gamma^d) = \bigotimes_{n=1}^d \ \mathcal{P}_{p_n}(\Gamma_n)$ with $\mathcal{P}_{p_n}(\Gamma_n) := \text{span}(y_n^m, \, m = 0, \ldots, p_n)$, $n = 1, \ldots, d$. Consider the univariate Lagrange interpolant along the n^{th} dimension of Γ^d where $\mathcal{I}_n^{m(i)} : C^0(\Gamma_n) \to \mathcal{P}_{m(i)-1}(\Gamma_n)$,

$$\mathcal{I}_n^{m(i)}(v(\mathbf{y})) := \sum_{k=1}^{m(i)} v(y_n^{(k)}) l_{n,k}(y_n^{(k)}),$$

$\{l_{n,k}\}_{k=1}^{m(i)}$ is a Lagrange basis of the space $\mathcal{P}_{p_n}(\Gamma_n)$, $i \geq 0$ is the level of approximation and $m(i) \in \mathbb{N}_+$ is the number of collocation nodes at level $i \in \mathbb{N}_+$ where $m(0) = 0$, $m(1) = 1$ and $m(i) \leq m(i+1)$ if $i \geq 1$.

We can now construct an interpolant by taking tensor products of $\mathcal{I}_n^{m(i)}$ along each dimension n. However, as d increases the dimension of $\mathcal{P}_p$ increases as $\prod_{n=1}^d (p_n + 1)$. Thus even for moderate d dimensions the computational cost of the Lagrange approximation becomes intractable. However, in the case of sufficient complex analytic regularity of the QoI with respect to the random variables defined on Γ^d, the application of Smolyak sparse grids are better suited. In the rest of this section the construction of the classical Smolyak sparse grid (see e.g. [6, 31]) is summarized. More details can be found in [4].

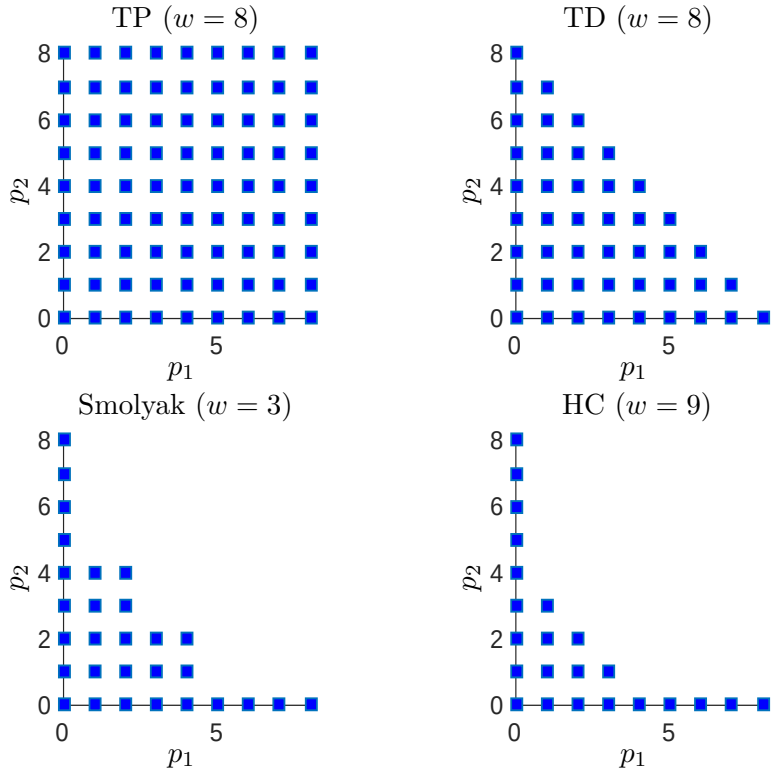

Figure 1: Index sets comparisons between Tensor Product, Total Degree, Smolyak and Hyperbolic Cross index sets for $d = 2$.

Consider the difference operator along the n^{th} dimension

$$\Delta_n^{m(i)} := \mathcal{I}_n^{m(i)} - \mathcal{I}_n^{m(i-1)}.$$

Given an integer $w \geq 0$, called the approximation level, and a multi-index $\mathbf{i} = (i_1, \ldots, i_d) \in \mathbb{N}_+^d$, let $g : \mathbb{N}_+^d \to \mathbb{N}$ be a strictly increasing function in each argument and define a sparse grid approximation of function $v(\mathbf{y}) \in C^0(\Gamma^d)$

$$\mathcal{S}_{w,d}^{m,g}[v(\mathbf{y})] = \sum_{\mathbf{i} \in \mathbb{N}_+^d : g(\mathbf{i}) \leq w} \bigotimes_{n=1}^{d} \Delta_n^{m(i_n)}(v(\mathbf{y})) \tag{6}$$

or equivalently written as

$$\mathcal{S}_{w,d}^{m,g}[v(\mathbf{y})] = \sum_{\mathbf{i} \in \mathbb{N}_+^d : g(\mathbf{i}) \leq w} c(\mathbf{i}) \bigotimes_{n=1}^{d} \mathcal{I}_n^{m(i_n)}(v(\mathbf{y})), \qquad \text{with } c(\mathbf{i}) = \sum_{\substack{\mathbf{j} \in \{0,1\}^d : \\ g(\mathbf{i}+\mathbf{j}) \leq w}} (-1)^{|\mathbf{j}|}. \tag{7}$$

From the previous expression, the sparse grid approximation is obtained as a linear combination of full tensor product interpolations. However, the constraint $g(\mathbf{i}) \leq w$ in (7) is typically chosen so as to forbid the use of tensor grids of high degree in all directions at the same time.

Let $\mathbf{m}(\mathbf{i}) = (m(i_1), \ldots, m(i_d))$ and consider the set of polynomial multi-degrees

$$\Lambda^{m,g}(w) = \{\mathbf{p} \in \mathbb{N}^d, \;\; g(\mathbf{m}^{-1}(\mathbf{p}+\mathbf{1})) \leq w\},$$

and corresponding set of monomials

$$\mathcal{Q}_{\Lambda^{m,g}(w)} := \left\{ \prod_{n=1}^{d} y_n^{p_n}, \;\; \text{with } \mathbf{p} \in \Lambda^{m,g}(w) \right\}.$$

Denote by $\mathbb{P}_{\Lambda^{m,g}(w)}(\Gamma^d)$ the corresponding multivariate polynomial space spanned by the monomials with multi-degree in $\Lambda^{m,g}(w)$, i.e.

$$\mathbb{P}_{\Lambda^{m,g}(w)}(\Gamma^d) = span \left\{ \prod_{n=1}^{d} y_n^{p_n}, \;\; \text{with } \mathbf{p} \in \Lambda^{m,g}(w) \right\}.$$

The following result proved in [4] states that the sparse approximation formula $\mathcal{S}_{w,d}^{m,g}$ is exact in $\mathbb{P}_{\Lambda^{m,g}(w)}(\Gamma_d)$:

Proposition 1.

a) For any $v \in C^0(\Gamma_d; V)$, we have $\mathcal{S}_{w,d}^{m,g}[v] \in \mathbb{P}_{\Lambda^{m,g}(w)} \otimes V$.

b) Moreover, $\mathcal{S}_{w,d}^{m,g}[v] = v, \;\; \forall v \in \mathbb{P}_{\Lambda^{m,g}(w)} \otimes V$.

Here V denotes a Banach space defined on Γ^d and

$$C^0(\Gamma_d; V) := \{v : \Gamma_d \to V \text{ is continuous on } \Gamma_d \text{ and } \max_{y \in \Gamma_d} \|v(y)\|_V < \infty\}.$$

The most typical choice of m and g is given by (see [6, 31])

$$m(i) = \begin{cases} 1, & \text{for } i = 1 \\ 2^{i-1}+1, & \text{for } i > 1 \end{cases} \quad \text{and} \quad g(\mathbf{i}) = \sum_{n=1}^{d} (i_n - 1).$$

Furthermore $\Lambda^{m.g}(w) := \{\mathbf{p} \in \mathbb{N}^d : \;\; \sum_n f(p_n) \leq f(w)\}$ where

$$f(p_n) = \begin{cases} 0, \; p_n = 0 \\ 1, \; p_n = 1 \\ \lceil \log_2(p_n) \rceil, \; p_n \geq 2 \end{cases}.$$

Other choices are shown in Table 2.

This choice of m and g combined with the choice of Clenshaw-Curtis (extrema of Chebyshev polynomials) abscissas leads to nested sequences of one dimensional interpolation formulas and a sparse grid with a highly reduced number of nodes compared to the corresponding tensor grid. Another good choice includes the Chebyshev Gauss-Lobatto abscissas. For any choice of $m(i) > 1$ the Clenshaw-Curtis abscissas are given by

$$y_j^n = -cos\left(\frac{\pi(j-1)}{m(i)-1}\right).$$

Approx. space	sparse grid: m, g	polynomial space: $\Lambda(w)$
Total Degree (TD)	$m(i) = i$ $g(\mathbf{i}) = \sum_n (i_n - 1) \leq w$	$\{\mathbf{p} \in \mathbb{N}^N : \ \sum_n p_n \leq w\}$
Hyperbolic Cross (HC)	$m(i) = i$ $g(\mathbf{i}) = \prod_n (i_n) \leq w + 1$	$\{\mathbf{p} \in \mathbb{N}^N : \ \prod_n (p_n + 1) \leq w + 1\}$
Smolyak (SM)	$m(i) = \begin{cases} 2^{i-1} + 1, \ i > 1 \\ 1, \ i = 1 \end{cases}$ $g(\mathbf{i}) = \sum_n (i_n - 1) \leq w$	$\{\mathbf{p} \in \mathbb{N}^N : \ \sum_n f(p_n) \leq w\}$ $f(p) = \begin{cases} 0, \ p = 0 \\ 1, \ p = 1 \\ \lceil \log_2(p) \rceil, \ p \geq 2 \end{cases}$

Table 2: Sparse grids approximations formulas for TD, HC and SM.

Suppose $u \in C^k_{mix}(\Gamma; \mathbb{R})$, which is the space of continuous mixed derivatives of degree k and is defined as

$$C^k_{mix}(\Gamma_d; \mathbb{R}) = \{v : \Gamma_d \to \mathbb{R} : \frac{\partial^{\alpha_1, \ldots, \alpha_d} v}{\partial^{\alpha_1} y_1 \ldots \partial^{\alpha_d} y_d} \in C^0(\Gamma_d; \mathbb{R}), \ n = 1, \ldots, d, \ \alpha_n \leq k\}.$$

The authors in [6] show that

$$\|v - \mathcal{S}^{m,g}_{w,d}[v]\|_{L^\infty(\Gamma)} \leq C(k, d) \|u\|_{C^k_{mix}(\Gamma)} w^{-k} (\log w)^{(k+2)(d-1)+1}),$$

where for all $v \in C^k_{mix}(\Gamma_d; \mathbb{R})$

$$\|v\|_{C^k_{mix}(\Gamma_d; \mathbb{R})} = \{v : \Gamma \to \mathbb{R} : \max_{\mathbf{y} \in \Gamma_d} \left| \frac{\partial^{\alpha_1, \ldots, \alpha_d} v(\mathbf{y})}{\partial^{\alpha_1} y_1 \ldots \partial^{\alpha_d} y_d} \right| < \infty\}.$$

The weakness of this bound is that the coefficient $C(k, d)$ is in general not known, thus making it difficult to estimate a-posteriori bounds on the solution. A better choice is the $L^\infty(\Gamma_d)$ bounds derived in [28]. These bounds are explicit, and the coefficients are directly based on analytic extensions of $v(\mathbf{y})$ on a well defined region $\Theta \subset \mathbb{C}^d$.

In [27, 28] the authors derive error estimates for isotropic and anisotropic Smolyak sparse grids with Clenshaw-Curtis and Gaussian abscissas, where the error $\|v - \mathcal{S}^{m,g}_{w,d}[v]\|_{L^\infty(\Gamma)}$ exhibits algebraic or sub-exponential convergence with respect to the number of collocation knots η (See Theorems 3.10, 3.11, 3.18 and 3.19 in [28] for more details). However, for these estimates to be valid $v \in C^0(\Gamma, \mathbb{R})$ has to admit an extension in the region defined by the polyellipse in $\mathbb{C}^d$, $\mathcal{E}_{\hat{\sigma}_1, \ldots, \hat{\sigma}_d} := \Pi^d_{n=1} \mathcal{E}_{n, \hat{\sigma}_n} \subset \Theta$, where

$$\mathcal{E}_{n, \hat{\sigma}_n} = \left\{ z \in \mathbb{C}, \hat{\sigma}_n \geq \delta_n \geq 0; \ \mathrm{Re}\, z = \frac{e^{\delta_n} + e^{-\delta_n}}{2} cos(\theta), \ \ \mathrm{Im}\, z = \frac{e^{\delta_n} - e^{-\delta_n}}{2} sin(\theta), \theta \in [0, 2\pi) \right\},$$

and $\hat{\sigma}_n > 0$ (see Figure 2).

Remark 1. *For simplicity, in [28], the error bounds assume that the term $M(v)$ is bounded by*

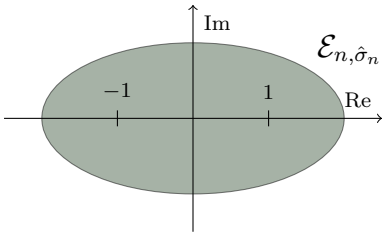

Figure 2: Polyellipse along the n^{th} dimension. The ellipse crosses the real axis at $\frac{e^{\hat{\sigma}_n}+e^{-\hat{\sigma}_n}}{2}$ and the imaginary axis at $\frac{e^{\hat{\sigma}_n}-e^{-\hat{\sigma}_n}}{2}$.

one. We remove this assumption and modify Theorem 3.10 and 3.11 in [28] to show the error bounds with the term $\tilde{M}(v)$, where

$$M(v) \leq \tilde{M}(v) = \max_{\mathbf{z} \in \mathcal{E}_{\hat{\sigma}_1,\ldots,\hat{\sigma}_d}} |v(\mathbf{z};\boldsymbol{\theta})|$$

We now introduce Chebyshev polynomials. These will be useful in deriving error estimates of the sparse grid. Let $T_k : \Gamma_1 \to \mathbb{R}$, $k = 0, 1, \ldots$, be a Chebyshev polynomial over $[-1, 1]$, which are defined recursively as:

$$T_0(y) = 1, T_1(y) = x$$
$$T_{k+1}(y) = 2xT_k(y) - T_{k-1}(y), \ldots$$

The following theorem describe the approximation of analytic functions with Chebyshev polynomials.

Theorem 1. *Let q be analytic and bounded by M on $\mathcal{E}_{\log\rho}$, $\rho > 1$, then the expansion*

$$q(y) = \alpha_0 + 2\sum_{k=1}^{\infty} \alpha_k T_k(y),$$

holds for all $y \in \mathcal{E}_{\log\rho}$ and $|\alpha_k| \leq M/\rho^k$. Furthermore if $y \in [-1, 1]$ then

$$|q(y) - \alpha_0 - 2\sum_{k=1}^{m} \alpha_k T_k(y)| \leq \frac{2M}{\rho - 1}\rho^{-m}$$

Since the interpolation operator $\mathcal{I}_n^{m(i)}$ is exact on the space $\mathcal{P}_{p_n-1}$, then if v is analytic on Γ_n we have from Theorem 1 that

$$\|(I - \mathcal{I}_n^{m(i)})v\|_{L^\infty(\Gamma_n)} \leq (1 + \Lambda_{m(i)})\frac{2M}{\rho - 1}\rho^{-m(i)},$$

where $\Lambda_{m(i)}$ is the Lebesgue constant and is bounded by $2\pi^{-1}(\log(m-1)+1)$ (see [3]). Thus, for $n = 1, \ldots, d$, with $\rho_n = \log\hat{\sigma}_n$ we have

$$\|(I - \mathcal{I}_n^{m(i)})v\|_{L^\infty(\Gamma^d)} \leq \tilde{M}(v)C(\sigma_n)ie^{-\sigma_n 2^i} \tag{8}$$

where $\sigma_n = \frac{\hat{\sigma}_n}{2} > 0$ and $C(\sigma_n) := \frac{2}{(e^{\sigma_n}-1)}$. We have that for all $n = 1, \ldots, d$

$$\begin{aligned}\|\Delta(v)^{m(i)}\|_{L^\infty(\mathcal{E}_{\hat{\sigma}_n})} &= \|(\mathcal{I}_n^{m(i)} - \mathcal{I}_n^{m(i-1)})v\|_{L^\infty(\Gamma^d)} \\ &\leq \|(I - \mathcal{I}_n^{m(i)})v\|_{L^\infty(\Gamma^d)} + \|(I - \mathcal{I}_n^{m(i-1)})v\|_{L^\infty(\Gamma^d)} \\ &\leq 2\tilde{M}(v)C(\sigma_n)ie^{-\sigma_n 2^{i-1}}\end{aligned} \tag{9}$$

and set

$$\sigma \equiv \min_{n=1,\ldots,d} \sigma_n,$$

i.e. for an isotropic sparse grid the overall asymptotic sub-exponential decay rate $\hat{\sigma}$ will be dominated by the smallest $\hat{\sigma}_n$.

Theorem 2. *Suppose that $v \in C^0(\Gamma; \mathbb{R})$ has an analytic extension on $\mathcal{E}_{\hat{\sigma}_1,\ldots,\hat{\sigma}_d}$ and bounded by $\tilde{M}(v)$. If $w > d/\log 2$ and a nested CC sparse grid is used then the following bound is valid:*

$$\|v - \mathcal{S}_{w,d}^{m,g} v\|_{L^\infty(\Gamma_d)} \leq \mathcal{Q}(\sigma, \delta^*(\sigma), d, \tilde{M}(v))\eta^{\mu_3(\sigma,\delta^*(\sigma),d)} \exp\left(-\frac{d\sigma}{2^{1/d}}\eta^{\mu_2(d)}\right), \tag{10}$$

where $\sigma = \hat{\sigma}/2$, $\delta^*(\sigma) := (e\log(2) - 1)/\tilde{C}_2(\sigma)$,

$$\mathcal{Q}(\sigma, \delta^*(\sigma), d, \tilde{M}(v)) := \frac{C_1(\sigma, \delta^*(\sigma), \tilde{M}(v))}{\exp(\sigma\delta^*(\sigma)\tilde{C}_2(\sigma))} \frac{\max\{1, C_1(\sigma, \delta^*(\sigma), \tilde{M}(v))\}^d}{|1 - C_1(\sigma, \delta^*(\sigma), \tilde{M}(v))|},$$

$\mu_2(d) = \frac{log(2)}{d(1+log(2d))}$ *and* $\mu_3(\sigma, \delta^*(\sigma), d) = \frac{\sigma\delta^*(\sigma)\tilde{C}_2(\sigma)}{1+\log(2d)}$. *Furthermore,* $C(\sigma) = \frac{4}{e^{2\sigma}-1}$,

$$\tilde{C}_2(\sigma) = 1 + \frac{1}{\log 2}\sqrt{\frac{\pi}{2\sigma}},\ \delta^*(\sigma) = \frac{e\log(2) - 1}{\tilde{C}_2(\sigma)},\ C_1(\sigma, \delta, \tilde{M}(v)) = \frac{4\tilde{M}(v)C(\sigma)a(\delta, \sigma)}{e\delta\sigma},$$

and

$$a(\delta, \sigma) := \exp\left(\delta\sigma\left\{\frac{1}{\sigma\log^2(2)} + \frac{1}{\log(2)\sqrt{2\sigma}} + 2\left(1 + \frac{1}{\log(2)}\sqrt{\frac{\pi}{2\sigma}}\right)\right\}\right).$$

Furthermore, if $w \leq d/\log 2$ then the following algebraic convergence bound is true:

$$\|v - \mathcal{S}_w^{m,g} v\|_{L^\infty(\Gamma_d)} \leq \frac{C_1(\sigma, \delta^*(\sigma), \tilde{M}(v))}{|1 - C_1(\sigma, \delta^*(\sigma), \tilde{M}(v))|} \max\left\{1, C_1(\sigma, \delta^*(\sigma), \tilde{M}(v))\right\}^d \eta^{-\mu_1}, \tag{11}$$

where $\mu_1 = \frac{\sigma}{1+\log(2d)}$.

Remark 2. *As the reader might have realized, it is not necessary to show that there is a bounded analytic extension on $\mathcal{E}_{\hat{\sigma}_1,\ldots,\hat{\sigma}_d}$. Since the convergence rate of the isotropic sparse grid is controlled by $\hat{\sigma}$, it is sufficient to show that v is bounded in $\mathcal{E}_{\hat{\sigma},\ldots,\hat{\sigma}} := \prod_{n=1}^d \mathcal{E}_{n,\hat{\sigma}}$ and the bound $\tilde{M}(v)$ will be smaller. However, an analytic extension on $\mathcal{E}_{\hat{\sigma}_1,\ldots,\hat{\sigma}_d}$ will be relevant for anisotropic sparse grids [27], which leads to faster convergence rates. We leave this as a future avenue to explore.*

Remark 3. *From Lemma 3.9 [28] the total number of collocation knots η satisfy the following bounds*

$$d(2^w - 1) \leq \eta \leq (2ed)^w \min\{w + 1, 2ed\}$$

Remark 4. *An alternative choice to the Smolyak sparse grid is the Hyperbolic Cross sparse grid. For this case we have that*

$$m(i) = i \quad \text{and} \quad g(\mathbf{i}) = \prod_{n=1}^{d} i_n$$

and $\Lambda^{m,g}(w) := \{\mathbf{p} \in \mathbb{N}^d : \ \prod_n (p_n + 1) \leq w\}$. *The cardinality of* $\Lambda^{m,g}(w)$ *for the HC grows much slower than SM or TD and works very well in practice. However, to my knowledge no* $L^\infty(\Gamma_d)$ *error bounds exist.*

Remark 5. *For many practical cases all the dimensions of* Γ^d *are not equally important. In this case the dimensionality of the sparse grid can be significantly reduced by means of a dimension adaptive [17] or anisotropic sparse grid. It is straightforward to build related anisotropic sparse approximation formulas by making the function* g *act differently on the input random variables* y_n. *Anisotropic sparse stochastic collocation [27] combines the advantages of isotropic sparse collocation with those of anisotropic full tensor product collocation.*

Another choice is to build quasi-optimal isotropic/anisotropic sparse grids [7, 26]. These grids have provable exponential convergence rates. However, in general, it is not shown how to synthesize quasi-optimal grids.

4 Multi-level approach

The general approach of this paper and multi-level basis construction are now presented. Denote by $\mathcal{P}^p(\mathbb{S})$ the span of the columns of the design matrix $\mathbf{M}$. Let $\mathbf{L} \in \mathbb{R}^{p \times N}$ be an orthogonal projection from $\mathbb{R}^n$ to $\mathcal{P}^p(\mathbb{S})$ and $\mathbf{W} \in \mathbb{R}^{(N-p)\times N}$ be an orthogonal projection from $\mathbb{R}^n$ to $\mathcal{P}^p(\mathbb{S})^\perp$ (the orthogonal complement of $\mathcal{P}^p(\mathbb{S})$). Moreover assume that the operator $\begin{bmatrix} \mathbf{W} \\ \mathbf{L} \end{bmatrix}$ is orthonormal.

By applying the operator $\mathbf{W}$ to (2) we obtain $\mathbf{Z}_W = \mathbf{W}\mathbf{Z} = \mathbf{W}(\mathbf{M}\boldsymbol{\beta} + \boldsymbol{\varepsilon}) = \mathbf{W}\boldsymbol{\varepsilon}$. Our first observation is that the trend contribution $\mathbf{M}\boldsymbol{\beta}$ is filtered out from the data $\mathbf{Z}$. We can now formulate the estimation of the covariance parameters $\boldsymbol{\theta}$ without the trend. The new log-likelihood function becomes

$$\ell_W(\boldsymbol{\theta}) = -\frac{n}{2}\log(2\pi) - \frac{1}{2}\log\det\{\mathbf{C}_W(\boldsymbol{\theta})\} - \frac{1}{2}\mathbf{Z}_W^{\mathrm{T}}\mathbf{C}_W(\boldsymbol{\theta})^{-1}\mathbf{Z}_W, \tag{12}$$

where $\mathbf{C}_W(\boldsymbol{\theta}) = \mathbf{W}\mathbf{C}(\boldsymbol{\theta})\mathbf{W}^{\mathrm{T}}$ and $\mathbf{Z}_W \sim \mathcal{N}_{N-p}(\mathbf{0}, \mathbf{W}\mathbf{C}(\boldsymbol{\theta})\mathbf{W}^{\mathrm{T}})$. A consequence of the filtering is that we obtain an unbiased estimator.

The decoupling of the likelihood function is not the only advantage of using $\mathbf{C}_{\mathbf{W}}(\boldsymbol{\theta})$. The following theorem also shows that $\mathbf{C}_{\mathbf{W}}(\boldsymbol{\theta})$ is more numerically stable than $\mathbf{C}(\boldsymbol{\theta})$.

Theorem 3. *Let* $\kappa(A) \to \mathbb{R}$ *be the condition number of the matrix* $A \in \mathbb{R}^{N\times N}$ *then*

$$\kappa(\mathbf{C}_{\mathbf{W}}(\boldsymbol{\theta})) \leq \kappa(\mathbf{C}(\boldsymbol{\theta})).$$

Proof. To see this let $\mathbf{v} := \mathbf{W}^{\mathrm{T}}\mathbf{w}$ for all $\mathbf{w} \in \mathbb{R}^{N-p}$, which implies that $\mathbf{v} \in \mathbb{R}^n \backslash \mathcal{P}^p(\mathbb{S})$. Moreover, this map is surjective. Now, $\mathbf{v}^{\mathrm{T}}\mathbf{C}(\boldsymbol{\theta})\mathbf{v} = \mathbf{w}^{\mathrm{T}}\mathbf{C}_{\mathbf{W}}(\boldsymbol{\theta})\mathbf{w}$ for all $\mathbf{w} \in \mathbb{R}^{N-p}$. Thus we have that for all $\mathbf{v} \in \mathbb{R}^n \backslash \mathcal{P}^p(\mathbb{S})$

$$\min_{\mathbf{v}\in\mathbb{R}^n\backslash\mathcal{P}^p(\mathbb{S})} \frac{\mathbf{v}^{\mathrm{T}}\mathbf{C}(\boldsymbol{\theta})\mathbf{v}}{\|\mathbf{v}\|^2} = \min_{\mathbf{w}\in\mathbb{R}^{N-p}} \frac{\mathbf{w}^{\mathrm{T}}\mathbf{C}_{\mathbf{W}}(\boldsymbol{\theta})\mathbf{w}}{\|\mathbf{w}\|^2} \quad \text{and} \quad \max_{\mathbf{v}\in\mathbb{R}^n\backslash\mathcal{P}^p(\mathbb{S})} \frac{\mathbf{v}^{\mathrm{T}}\mathbf{C}(\boldsymbol{\theta})\mathbf{v}}{\|\mathbf{v}\|^2} = \max_{\mathbf{w}\in\mathbb{R}^{N-p}} \frac{\mathbf{w}^{\mathrm{T}}\mathbf{C}_{\mathbf{W}}(\boldsymbol{\theta})\mathbf{w}}{\|\mathbf{w}\|^2}.$$

Now, it is not hard to see that

$$0 < \min_{\mathbf{v}\in\mathbb{R}^n} \frac{\mathbf{v}^{\mathrm{T}}\mathbf{C}(\boldsymbol{\theta})\mathbf{v}}{\|\mathbf{v}\|^2} \leq \min_{\mathbf{v}\in\mathbb{R}^n\setminus\mathcal{P}^p(\mathbb{S})} \frac{\mathbf{v}^{\mathrm{T}}\mathbf{C}(\boldsymbol{\theta})\mathbf{v}}{\|\mathbf{v}\|^2} \leq \max_{\mathbf{v}\in\mathbb{R}^n\setminus\mathcal{P}^p(\mathbb{S})} \frac{\mathbf{v}^{\mathrm{T}}\mathbf{C}(\boldsymbol{\theta})\mathbf{v}}{\|\mathbf{v}\|^2} \leq \max_{\mathbf{v}\in\mathbb{R}^n} \frac{\mathbf{v}^{\mathrm{T}}\mathbf{C}(\boldsymbol{\theta})\mathbf{v}}{\|\mathbf{v}\|^2}.$$

□

Theorem 3 states that the condition number of $\mathbf{C}_W(\boldsymbol{\theta})$ is less or equal to the condition number of $\mathbf{C}(\boldsymbol{\theta})$. Thus computing the inverse of $\mathbf{C}_W(\boldsymbol{\theta})$ (using a direct or iterative method) will generally be more stable.

In practice, computing the inverse of $\mathbf{C}_W(\boldsymbol{\theta})$ will be much more stable than $\mathbf{C}(\boldsymbol{\theta})$ depending on the choice of the index set $\Lambda(w)$. This has many significant implications as it will now be possible to solve numerically unstable problems.

There are other advantages to the structure of the matrix $\mathbf{C}_W(\boldsymbol{\theta})$. In section 7 we show that for a good choice of the $\mathcal{P}(\mathbb{S})$ the entries of $\mathbf{C}_W(\boldsymbol{\theta})$ decay rapidly, and most of the entries can be safely eliminated without loosing accuracy. A level dependent criterion approach is shown in section 5 that indicate which entries are computed and which ones are not. With this approach a sparse covariance matrix $\tilde{\mathbf{C}}_W$ is constructed such that it is close to $\mathbf{C}_W$ in a matrix norm sense, even if the observations are highly correlated with distance.

4.1 Random projection multi-Level basis

In this section the construction of Multi-Level Basis (MB) is shown. The MB can then be used to: (i) form the multi-level likelihood (12), (ii) sparsify the covariance matrix $\mathbf{C}_W(\boldsymbol{\theta})$, and (iii) improve the conditioning over the covariance matrix $\mathbf{C}(\boldsymbol{\theta})$. But first, let us establish notations and definitions:

- Given $\mathcal{Q}_{\Lambda^{m,g}(w)}$ and the locations $\mathbb{S}$ construct the design matrix $\mathbf{M}$. Furthermore, form a second set of monomials $\tilde{\mathcal{Q}}^a_{\Lambda^{m,g}(w)} := \mathcal{Q}_{\Lambda^{m,g}(w+a)}$ for $a = 0, 1, \ldots$, i.e. $\mathcal{Q}_{\Lambda^{m,g}(w)} \subset \tilde{\mathcal{Q}}^a_{\Lambda^{m,g}(w)}$. Denote the accuracy parameter $\tilde{p} \in \mathbb{N}$ as the cardinality of $\tilde{\mathcal{Q}}^a_{\Lambda^{m,g}(w)}$. From the set of monomials $\tilde{\mathcal{Q}}^a_{\Lambda^{m,g}(w)}$, for some user given parameter $a \in \mathbb{N}_0$, and the set of observations $\mathbb{S}$ generate the design matrix $\tilde{\mathbf{M}}^a$. Denote also the space $\mathcal{P}^{\tilde{p}}(\mathbb{S})$ as the span of the columns of $\tilde{\mathbf{M}}^a$.

- For any index $i, j \in \mathbb{N}_0$, $1 \leq i \leq N$, $1 \leq j \leq N$, let $\mathbf{e}_i[j] = \delta[i-j]$, where $\delta[\cdot]$ is the discrete Kronecker delta function.

- Let $\phi(\mathbf{x}, \mathbf{y}; \boldsymbol{\theta}) : \Gamma_d \times \Gamma_d \to \mathbb{R}$ be the covariance function and assumed to be a positive definite. Let $\mathbf{C}(\boldsymbol{\theta})$ be the covariance matrix that is formed from all the interactions between the observation locations $\mathbb{S}$ i.e. $\mathbf{C}_W(\boldsymbol{\theta}) := \{\phi(\mathbf{x}_i, \mathbf{y}_j)\}$, where $i, j, = 1, \ldots, N$. Alternatively we refer to $\phi(r; \boldsymbol{\theta})$ as the covariance function where $r := \|\mathbf{x} - \mathbf{y}\|_2$.

Definition 1. *The Matérn covariance function:*

$$\phi(r; \boldsymbol{\theta}) = \frac{1}{\Gamma(\nu)2^{\nu-1}} \left(\sqrt{2\nu}\frac{r}{\rho}\right)^{\nu} K_\nu \left(\sqrt{2\nu}\frac{r}{\rho}\right),$$

where Γ is the gamma function, $0 < \nu$, $0 < \rho < \infty$, and K_ν is the modified Bessel function of the second kind.

Remark 6. *The Matérn covariance function is a good choice for the random field model. The parameter ρ controls the length correlation and the parameter ν changes the shape. For example, if $\nu = 1/2 + n$, where $n \in \mathbb{N}_+$, then (see [2])*

$$\phi(r;\rho) = \exp\left(-\frac{\sqrt{2\nu}r}{\rho}\right)\frac{\Gamma(n+1)}{\Gamma(2n+1)}\sum_{k=1}^{n}\frac{(n+1)!}{k!(n-k)!}\left(\frac{\sqrt{8\nu}r}{\rho}\right)^{n-k}$$

and for $\nu \to \infty$

$$\nu \to \infty \Rightarrow \phi(r;\boldsymbol{\theta}) \to \exp\left(-\frac{r^2}{2\rho^2}\right).$$

The first step is to decompose the domain Γ^d into a multi-level domain decomposition. A first choice that comes to mind is a kD-tree decomposition of the space $\mathbb{R}^d$. This is a good choice for lower dimensions, however, as the number of dimensions d become larger a better approach is to use a Random Projection (RP) tree [12].

First start with the root node B_0^0 at level 0 that contains all the observation nodes in $\mathbb{S}$. Now, split these nodes into two children cells B_1^1 and B_2^1 at level 1 according to the following rule:

1. Generate unit random vector v in $\mathbb{R}^d$
2. Project all the nodes $\mathbf{x} \in \mathbb{S}$ in the cell onto the unit vector v.
3. Split the cell with respect to the median of the projections.

For each cell B_1^1 and B_2^1 repeat the procedure until there is at most p nodes at the leaf nodes. Thus a binary tree is obtained, which is of the form B_0^0, B_1^1, B_2^1, B_3^2, B_4^2, B_5^2, B_6^2, $\ldots$, where t is the maximal depth (level) of the tree. Now, let $\mathcal{B}$ be the set of all the cells in the tree and $\mathcal{B}^n$ be the set of all the cells at level $0 \leq n \leq t$. The construction of the RP tree is described in Algorithms 1 and 2. In addition, for each cell a unique node number, current tree depth, threshold level and projection vector are also assigned. This will be useful for searching the tree.

Remark 7. *A kD-tree can also be constructed with Algorithms 1 and 3.*

Now, suppose there is a one-to-one mapping between the set of unit vectors $\mathcal{E} := \{\mathbf{e}_1, \ldots, \mathbf{e}_N\}$, which is denoted as leaf unit vectors, and the set of locations $\{\mathbf{x}_1, \ldots, \mathbf{x}_N\}$, i.e. $\mathbf{x}_n \longleftrightarrow \mathbf{e}_n$ for all $n = 1, \ldots, N$. It is clear that the span of the vectors $\{\mathbf{e}_1, \ldots, \mathbf{e}_N\}$ is $\mathbb{R}^N$. The next step is to construct a new basis of $\mathbb{R}^n$ that is multi-level and orthonormal.

(a) Start at the maximum level of the random projection tree, i.e. $q = t$.

(b) For each leaf cell $B_k^q \in \mathcal{B}^q$ assume without loss of generality that there are s observations nodes $\mathbb{S}^q := \{\mathbf{x}_1, \ldots, \mathbf{x}_s\}$ with associated vectors $C_k^q := \{\mathbf{e}_1, \ldots, \mathbf{e}_s\}$. Denote $\mathcal{C}_k^q$ as the span of the vectors in C_k^q.

   i) Let $\boldsymbol{\phi}_j^q := \sum_{\mathbf{e}_i \in B_k^q} c_{i,j}^q \mathbf{e}_i$, $j = 1, \ldots, a$; $\quad \boldsymbol{\psi}_j^q := \sum_{\mathbf{e}_i \in B_k^q} d_{i,j}^q \mathbf{e}_i$, $j = a+1, \ldots, s$, where $c_{i,j}^q$, $d_{i,j}^q \in \mathbb{R}$ and for some $a \in \mathbb{N}^+$. It is desired that the new discrete MB vector $\boldsymbol{\psi}_j^q$ be orthogonal to $\mathcal{P}^{\tilde{p}}(\mathbb{S})$, i.e., for all $g \in \mathcal{P}^{\tilde{p}}(\mathbb{S})$:

$$\sum_{i=1}^{n} g[i]\boldsymbol{\psi}_j^q[i] = 0 \tag{13}$$

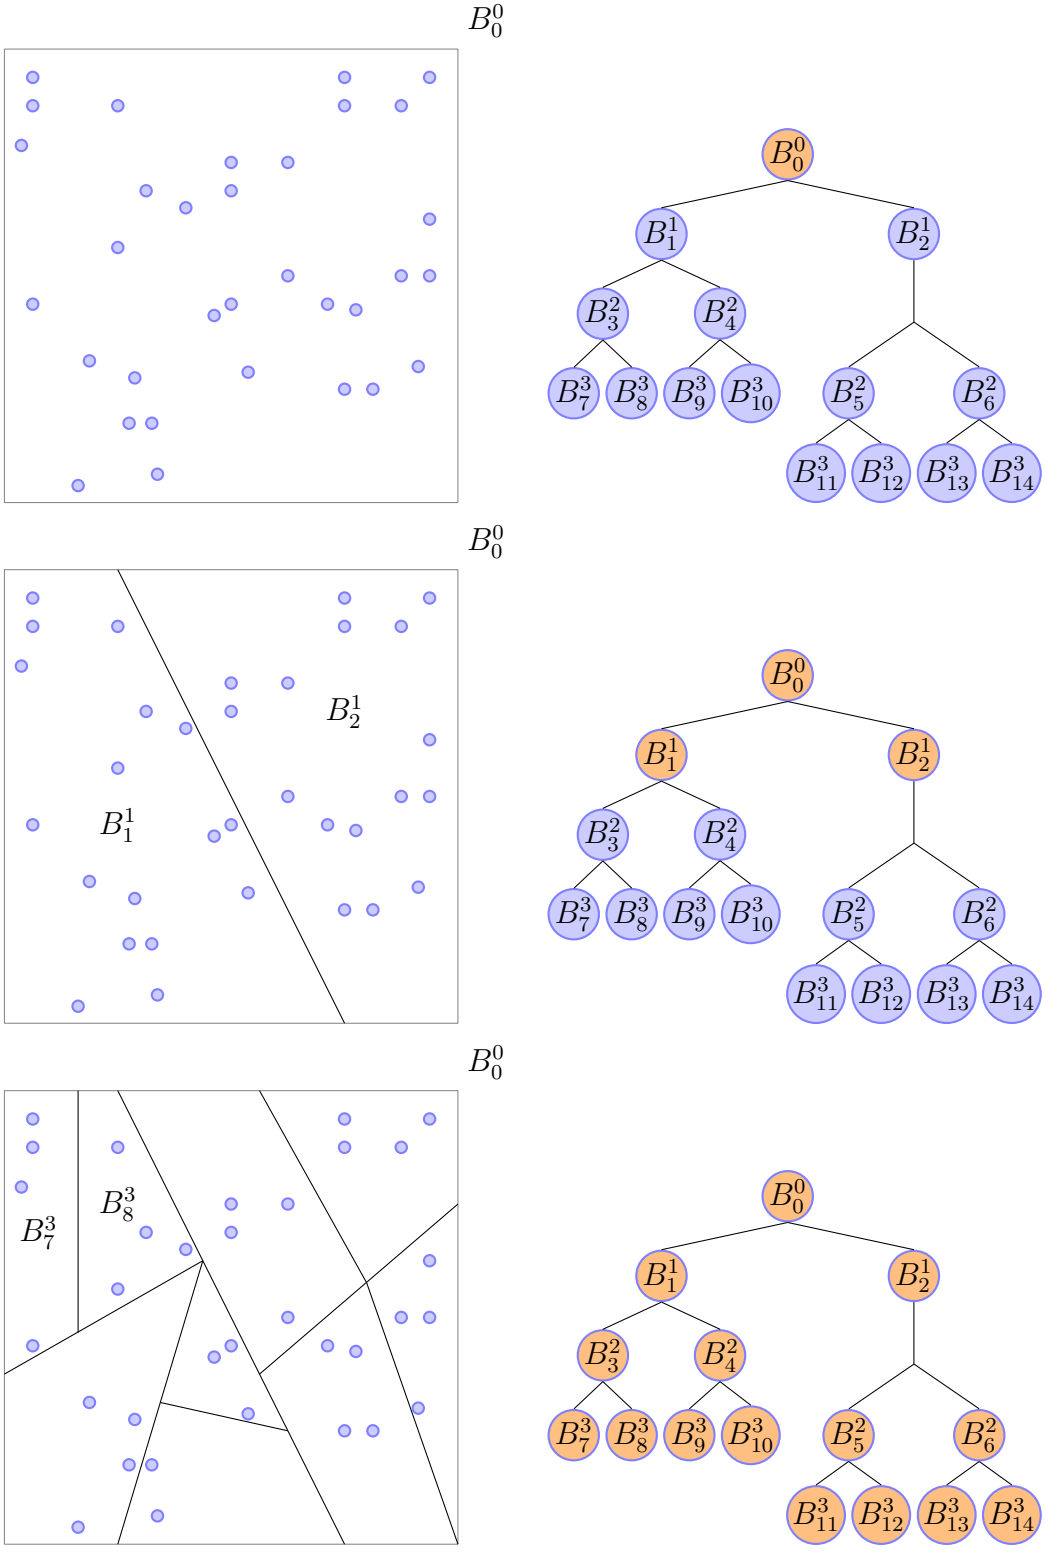

Figure 3: Random projection tree construction.

Input: $\mathbb{S}$, node, currentdepth, n_0
Output: Tree, node
begin
  if *Tree = root* **then**
    node $\leftarrow$ 0, currentdepth $\leftarrow$ 0 Tree $\leftarrow$ MakeTree($\mathbb{S}$, node, currentdepth + 1, n_0)
  else
    Tree.node = node
    Tree.currentdepth = currentdepth - 1
    node $\leftarrow$ node + 1
    if $|\mathbb{S}| < n_0$ **then**
      return (Leaf)
    end
    (Rule, threshold, v) $\leftarrow$ ChooseRule($\mathbb{S}$)
    (Tree.LeftTree, node) $\leftarrow$ MakeTree($\mathbf{x} \in \mathbb{S}$: Rule($\mathbf{x}$) = True, node, currentdepth + 1, n_0)
    (Tree.RightTree, node) $\leftarrow$ MakeTree($\mathbf{x} \in \mathbb{S}$: Rule($\mathbf{x}$) = false, node, currentdepth + 1, n_0)
    Tree.threshold = threshold
    Tree.v = v
  end
end

Algorithm 1: MakeTree($\mathbb{S}$) function

Input: $\mathbb{S}$
Output: Rule, threshold, v
begin
  choose a random unit vector v
  Rule(x) := $x \cdot v \leq$ threshold = median $\{z \cdot v : z \in \mathbb{S}\}$
end

Algorithm 2: ChooseRule($\mathbb{S}$) function for RP tree

ii) Form the matrix $\mathcal{M}^q := (\tilde{\mathbf{M}}^a)^{\mathrm{T}}\mathbf{V}$, where $\mathcal{M}^q \in \mathbb{R}^{p\times\mathcal{S}}$, $\mathbf{V} \in \mathbb{R}^{n\times s}$, and $\mathbf{V} := [\ldots, \mathbf{e}_i, \ldots]$ for all $\mathbf{e}_i \in C_k^q$. Now, suppose that the matrix $\mathcal{M}^q$ has rank a and then perform the Singular Value Decomposition (SVD). Denote by $\mathbf{UDV}$ the SVD of $\mathcal{M}^q$, where $\mathbf{U} \in \mathbb{R}^{\tilde{p}\times s}$, $\mathbf{D} \in \mathbb{R}^{s\times s}$, and $\mathbf{V} \in \mathbb{R}^{s\times s}$.

iii) To satisfy equation (13) make the following choice

$$\left[\begin{array}{ccc|ccc} c_{0,1}^q & \cdots & c_{a,1}^q & d_{a+1,1}^q & \cdots & d_{s,1}^q \\ c_{0,2}^q & \cdots & c_{a,2}^q & d_{a+1,2}^q & \cdots & d_{s,2}^q \\ \vdots & \vdots & \vdots & \vdots & \vdots & \vdots \\ c_{0,s}^q & \cdots & c_{a,s}^q & d_{a+1,s}^q & \cdots & d_{s,s}^q \end{array}\right] := \mathbf{V}^{\mathrm{T}},$$

where the columns $a+1, \ldots, s$ form an orthonormal basis of the nullspace $N_0(\mathcal{M})$. Similarly, the columns $1, \ldots, a$ form an orthonormal basis of $\mathbb{R}^s \backslash N_0(\mathcal{M})$. Since the vectors in C_k^q are orthonormal then $\boldsymbol{\phi}_1^q, \ldots, \boldsymbol{\phi}_a^q, \boldsymbol{\psi}_{a+1}^q, \ldots, \boldsymbol{\psi}^q$ form an orthonormal basis of $\mathcal{S}$. Moreover $\boldsymbol{\psi}_{a+1}^q, \ldots, \boldsymbol{\psi}^q$ satisfy equation (13), i.e., are orthogonal to $\mathcal{P}^{\tilde{p}}(\mathbb{S})$ and are locally adapted to the locations contained in the cell B_k^q.

Input: $\mathbb{S}$
Output: Rule, threshold, v
begin
  choose a coordinate direction
  Rule(x) := $x \cdot v \leq$ threshold = median
end

Algorithm 3: ChooseRule($\mathbb{S}$) function for kD-tree

iv) Denote by D_k^q the collection of all the vectors $\boldsymbol{\psi}_{a+1}^q, \ldots, \boldsymbol{\psi}^q$. Notice that the vectors $\boldsymbol{\phi}_1^q, \ldots, \boldsymbol{\phi}_a^q$, which are denoted with a slight abuse of notation as the scaling vectors, are *not* orthogonal to $\mathcal{P}^{\tilde{p}}(\mathbb{S})$. They need to be further processed.

v) Let $\mathcal{D}^q$ be the union of the vectors in D_k^q for all the cells $B_k^q \in \mathcal{B}_k^q$. Denote by $W_q(\mathbb{S})$ as the span of all the vectors in $\mathcal{D}^q$.

(c) The next step is to go to level $q-1$. For any two sibling cells denote B_{left}^q and B_{right}^q at level q denote $C_{\tilde{k}}^{q-1}$ as the collection of the scaling functions from both cells, for some index $\tilde{k}$.

(d) Let $q := q-1$. If $B_k^q \in \mathcal{B}^q$ is a leaf cell then repeat steps (b) to (d). However, if $B_k^q \in \mathcal{B}^q$ is not a leaf cell, then repeat steps (b) to (d), but replace the leaf unit vectors with the scaling vectors contained in C_k^q with $C_{\tilde{k}}^{q-1}$.

(e) When $q = -1$ is reached repeat steps (b) to (d), but replace $\tilde{p}$ with p, e.g. $\mathcal{P}^{\tilde{p}}(\mathbb{S})$ with $\mathcal{P}^p(\mathbb{S})$. The ML basis vectors will span the space $W_{-1}(\mathbb{S}) := \mathcal{P}^{\tilde{p}}(\mathbb{S}) \backslash \mathcal{P}^p(\mathbb{S})$.

When the algorithm stops a series orthogonal subspaces $W_{-1}(\mathbb{S}), \ldots, W_t(\mathbb{S})$ (and their corresponding basis vectors) are obtained. These subspaces are orthogonal to $V_{-1}(\mathbb{S}) := span\{\phi_1^{-1}, \ldots, \phi_p^{-1}\}$. Note that the orthonormal basis vectors of $V_{-1}(\mathbb{S})$ also span the space $\mathcal{P}^p(\mathbb{S})$.

Remark 8. *Following Lemma 2 in [9] it can be shown that*

$$\mathbb{R}^N = \mathcal{P}^p(\mathbb{S}) \oplus W_{-1}(\mathbb{S}) \oplus W_0(\mathbb{S}) \oplus W_1(\mathbb{S}) \oplus \cdots \oplus W_t(\mathbb{S}),$$

where $W_{-1}(\mathbb{S}) := \mathcal{P}^{\tilde{p}}(\mathbb{S}) \backslash \mathcal{P}^p(\mathbb{S})$. *Also, it can then be shown that at most* $\mathcal{O}(Nt)$ *computational steps are needed to construct the multi-level basis of* $\mathbb{R}^N$.

From the basis vectors of the subspaces $(\mathcal{P}^p(\mathbb{S}))^\perp = \cup_{i=-1}^t W_i(\mathbb{S})$ an orthogonal projection matrix $\mathbf{W} : \mathbb{R}^N \to (\mathcal{P}^p(\mathbb{S}))^\perp$ can be built. The dimensions of $\mathbf{W}$ is $(N-p) \times N$ since the total number of orthonormal vectors that span $\mathcal{P}^p(\mathbb{S})$ is p. Conversely, the total number of orthonormal vectors that span $(\mathcal{P}^p(\mathbb{S}))^\perp$ is $N-p$.

Let $\mathbf{L}$ be a matrix where each row is an orthonormal basis vector of $\mathcal{P}^p(\mathbb{S})$. For $i = -1, \ldots, t$ let $\mathbf{W}_i$ be a matrix where each row is a basis vector of the space $W_i(\mathbb{S})$. The matrix $\mathbf{W} \in \mathbb{R}^{(N-p)\times N}$ can now be formed, where $\mathbf{W} := \left[\mathbf{W}_t^{\mathrm{T}}, \ldots, \mathbf{W}_0^{\mathrm{T}}, \mathbf{W}_{-1}^{\mathrm{T}}\right]^{\mathrm{T}}$, therefore a) The matrix $\mathbf{P} := \begin{bmatrix} \mathbf{W} \\ \mathbf{L} \end{bmatrix}$ is orthonormal, i.e., $\mathbf{P}\mathbf{P}^{\mathrm{T}} = \mathbf{I}$; b) Any vector $\mathbf{v} \in \mathbb{R}^n$ can be written as $\mathbf{v} = \mathbf{L}^{\mathrm{T}}\mathbf{v}_L + \mathbf{W}^{\mathrm{T}}\mathbf{v}_W$ where $\mathbf{v}_L \in \mathbb{R}^p$ and $\mathbf{v}_W \in \mathbb{R}^{N-p}$ are unique; c) The matrix $\mathbf{W}$ contains at most $\mathcal{O}(Nt)$ non-zero entries and $\mathbf{L}$ contains at most $\mathcal{O}(Np)$ non-zero entries. This implies that for any vector $\mathbf{v} \in \mathbb{R}^n$ the computational cost of applying $\mathbf{W}\mathbf{v}$ is at most $\mathcal{O}(Nt)$ and $\mathbf{L}\mathbf{v}$ is at most $\mathcal{O}(Np)$.

Lemma 1. *Assuming that $n_0 = 2\tilde{p}$, for any level $q = 0, \ldots, t$ there is at most $\tilde{p}2^q$ multi-level basis vectors. For level $q = -1$ there is at most $p - \tilde{p}$ multi-level vectors.*

Proof. By construction any leaf cell has at most $\tilde{p}$ multi-level vectors and $\tilde{p}$ scaling vectors that are to be used for the next level.

Starting at the finest level t, for each cell $B_k^t \in \mathcal{B}^t$ there is at most $\tilde{p}$ multi-level vectors and $\tilde{p}$ scaling vectors that are to be used for the next level. Since there is at most 2^t cells then there is at most $2^t\tilde{p}$ multi-level vectors.

Now, for each pair of left and right (siblings) cells at level t the parent cell at level $t-1$ will have at most $2\tilde{p}$ scaling functions. Thus at most $\tilde{p}$ multi-level vectors and $\tilde{p}$ scaling vectors are obtained that are to be used for the next level. Now, the rest of the cells at level t are leafs and will have at most $\tilde{p}$ multi-level vectors and $\tilde{p}$ scaling vectors that are to be used for the next level. Since there is at most 2^{t-1} cells at level $t-1$, there is at most $2^{t-1}\tilde{p}$ multi-level vectors. Now, follow an inductive argument until $q = 0$. Finally at level $q = -1$ there will be at most $p - \tilde{p}$ □

Lemma 2. *For any level $q = 0, \ldots, t$ any multi-level vector $\boldsymbol{\psi}_m^q$ associated with a cell $B_k^q \in \mathcal{B}^q$ has at most $2^{t-q+1}\tilde{p}$ non zero entries.*

Proof. For any leaf cell at the bottom of the tree (level t) there is at most $2\tilde{p}$ observations, thus the number of non zero entries of level q multi-level vectors is $2\tilde{p}$. Combining the left and right cells, the parent cell has at most $4\tilde{p}$ observations, thus the associated multi-level vectors has $4\tilde{p}$ non zero entries. By induction at any level l the number of nonzero entries is at most $2^{t-q+1}\tilde{p}$.

Now for any leaf cell at any other level $l < t$ the number of nonzero entries is at most $2\tilde{p}$. Following an inductive argument the result is obtained. □

Remark 9. Extended multi-level basis: *A variation to the multi-level basis is obtained by replacing $\tilde{\mathcal{Q}}^a_{\Lambda^{m,g}(w)}$, $a \in \mathbb{N}_+$, with an extended set of monomials. Let $\tilde{\Lambda}^{m,g}(w)$, $w \in \mathbb{N}_+$, be an index set such that*

$$\tilde{\Lambda}^{m,g}(w) := \{\mathbf{p} \in \mathbb{N}_+^d, 2\mathbf{p} \in \mathbb{N}_+^d : \mathbf{p} \in \Lambda^{m,g}(w)\},$$

and any repeated index in $\tilde{\Lambda}^{m,g}(w)$ is removed. From this index set form the set of monomials $\tilde{\mathcal{Q}}^a_{\tilde{\Lambda}^{m,g}(w)}$, for all $a \in \mathbb{N}_+$. As defined before, the accuracy parameter $\tilde{p} \in \mathbb{N}$ is the cardinality of $\tilde{\mathcal{Q}}^a_{\tilde{\Lambda}^{m,g}(w)}$ and the design matrix $\tilde{\mathbf{M}}^a$ is formed from $\tilde{\mathcal{Q}}^a_{\tilde{\Lambda}^{m,g}(w)}$ and $\mathbb{S}$.

5 Multi-level covariance matrix

The multi-level covariance matrix $\mathbf{C}_W(\boldsymbol{\theta})$ and sparse version $\tilde{\mathbf{C}}_W(\boldsymbol{\theta})$ can be now constructed. Recall from the discussion in Section 4 that $\mathbf{C}_W(\boldsymbol{\theta}) := \mathbf{W}\mathbf{C}(\boldsymbol{\theta})\mathbf{W}^{\mathrm{T}}$. From the multi-level basis construct in Section 4.1 the following operator is built: $\mathbf{W} := \left[\mathbf{W}_t^{\mathrm{T}}, \ldots, \mathbf{W}_0^{\mathrm{T}}, \mathbf{W}_{-1}^{\mathrm{T}}\right]^{\mathrm{T}}$. Thus the covariance matrix $\mathbf{C}(\boldsymbol{\theta})$ is transformed into $\mathbf{C}_W(\boldsymbol{\theta})$, where each of the blocks $\mathbf{C}_W^{i,j}(\boldsymbol{\theta}) = \mathbf{W}_i\mathbf{C}(\boldsymbol{\theta})\mathbf{W}_j^{\mathrm{T}}$ are formed from all the interactions of the MB vectors between levels i and j, for all $i, j = -1, \ldots, t$. The structure of $\mathbf{C}_W(\boldsymbol{\theta})$ is shown in Figure 4. Thus for any $\boldsymbol{\psi}_{\tilde{l}}^i$ and $\boldsymbol{\psi}_{\tilde{k}}^j$ vectors there is a unique entry of $\mathbf{C}_W^{i,j}$ of the form $(\boldsymbol{\psi}_{\tilde{k}}^i)^{\mathrm{T}}\mathbf{C}(\boldsymbol{\theta})\boldsymbol{\psi}_{\tilde{l}}^j$. The blocks $\mathbf{C}_W^{i,j}$, where $i = -1$ or $j = -1$, correspond to the case where the accuracy term $\tilde{p} > p$.

5.1 Sparsification of multi-level covariance matrix

A sparse version of the covariance matrix $\mathbf{C}_W(\boldsymbol{\theta})$ can be built by using a level and distance dependent strategy:

1. Given a cell B_k^l at level $l \geq -1$ identify the corresponding tree node value Tree.node and the tree depth Tree.currentdepth. Note that the Tree.currentdepth and the MB level q are the same for $q = 0, \ldots, t$. However, for $q = -1$ the MB is associated to the Tree.currentdepth $= 0$.

2. Let $\mathbb{K} \subset \mathbb{S}$ be all the observations nodes contained in the cell B_k^l.

3. Let $\tau \geq 0$ be the distance parameter given by the user.

4. Let the Targetdepth be equal to the desired level of the tree. In the case that it is -1 then the Targetdepth is zero.

The objective now is to find all the cells at the Targetdepth that overlap a ball with radius τ from the projection of each of the nodes $\mathbf{x} \in \mathbb{K}$ onto the vector Tree.v. This is done by searching the tree from the root node. At each traversed node check that all the nodes $\mathbf{x} \in \mathbb{K}$ satisfy the following rule:

$$\mathbf{x} \cdot \text{Tree}.v + \tau \leq Tree.threshold.$$

If this is true then the search continues down the left tree, otherwise both the left and right tree are searched. The full search algorithm is described in Algorithms 4, 5, and 6.

Input: Tree, $\mathbb{K}$, Targetdepth, τ
Output: Targetnodes
begin
  Targetnodes $\leftarrow \emptyset$
  Targetnodes $\leftarrow$ LocalSearchTree(Tree, $\mathbb{K}$, Targetdepth, τ, Targetnodes);
end

Algorithm 4: SearchTree function(Tree, $\mathbb{K}$, Targetdepth, τ)

Input: Tree, $\mathbb{K}$, Targetdepth, τ, Targetnodes
Output: Targetnodes
begin
  if *Targetdepth = Tree.currentdepth* **then**
    return Targetnodes = Targetnodes $\cup$ Tree.node
  end
  if *Tree = leaf* **then**
    return
  end
  Rule $\leftarrow$ ChooseSecondRule($\mathbb{K}$, Tree, τ)
  if *Rule(*$\mathbf{x}$*)=true* $\forall \mathbf{x} \in \mathbb{K}$ **then**
    Targetnodes $\leftarrow$ LocalSearchTree(Tree.LeftTree, $\mathbb{K}$, Targetdepth, τ, Targetnodes)
  else
    Targetnodes $\leftarrow$ LocalSearchTree(Tree.LeftTree, $\mathbb{K}$, Targetdepth, τ, Targetnodes)
    Targetnodes $\leftarrow$ LocalSearchTree(Tree.RightTree, $\mathbb{K}$, Targetdepth, τ, Targetnodes)
  end
end

Algorithm 5: LocalSearchTree(Tree, $\mathbb{K}$, Targetdepth, τ) function

Input: $\mathbb{K}$, Tree, τ
Output: Rule
begin
  Rule($\mathbf{x}$) := $\mathbf{x} \cdot$ Tree.$v + \tau \leq$ Tree.threshold
end

Algorithm 6: ChooseSecondRule($\mathbb{K}$) function

The sparse matrix blocks $\mathbf{C}_W^{i,j}(\boldsymbol{\theta})$ can be built with a simple rule from Algorithm 6. Compute all the entries of $\mathbf{C}_W^{i,j}(\boldsymbol{\theta})$ that correspond to the interactions between any two cells $B_k^i \in \mathcal{B}^i$ and $B_l^j \in \mathcal{B}^j$ such that the maximum Euclidean distance between any two observations $\mathbf{x}^* \in B_k^i$ and $\mathbf{y}^* \in B_l^j$ is less or equal to $\tau_{i,j} \geq 0$ i.e. $\|\mathbf{x}^* - \mathbf{y}^*\| \leq \tau_{i,j}$. Thus compute the entries $(\boldsymbol{\psi}_{\tilde{k}}^i)^{\mathrm{T}} \mathbf{C}(\boldsymbol{\theta}) \boldsymbol{\psi}_{\tilde{l}}^j$ for all $\psi_{\tilde{k}}^i \in D^i$ (i.e. columns of D^i) and $\psi_{\tilde{l}}^j \in D^j$ (see Algorithm 7).

In Figure 5 an example for searching local neighborhood cells of randomly placed observations in $\mathbb{R}^2$ is shown. The orange nodes correspond to the source cell. By choosing a suitable value for τ the blue nodes in the immediate cell neighborhood are found by using Algorithms 4, 5 and 6.

Remark 10. *Since the matrix $\tilde{\mathbf{C}}_W(\boldsymbol{\theta})$ is symmetrical it is only necessary to compute the blocks $\mathbf{C}_W^{i,j}(\boldsymbol{\theta})$ for $i = 1, \ldots, t$ and $j = i, \ldots t$.*

Input: Tree, i, j, $\tau_{i,j}$, $\mathcal{B}^i$, $\mathcal{B}^j$, $\mathcal{D}^i$, $\mathcal{D}^i$, $\mathbf{C}(\boldsymbol{\theta})$
Output: $\tilde{\mathbf{C}}_W^{i,j}(\boldsymbol{\theta})$
begin
  Targetnodes $\leftarrow \emptyset$
  for $B_m^i \in \mathcal{B}^i$ **do**
    $\mathbb{K} \leftarrow B_m^i$
    for $B_q^j \leftarrow$ *LocalSearchTree(Tree,* $\mathbb{K}$*, Targetdepth* (i)*,* $\tau_{i,j}$*, Targetnodes)* **do**
      for $\psi_k^i \in D^i$ **do**
        for $\psi_l^j \in D^j$ **do**
          Compute $(\boldsymbol{\psi}_k^i)^{\mathrm{T}} \mathbf{C}(\boldsymbol{\theta}) \boldsymbol{\psi}_l^j$ in $\tilde{\mathbf{C}}_W^{i,j}(\boldsymbol{\theta})$
        end
      end
    end
  end
end

Algorithm 7: Construction of sparse matrix $\tilde{\mathbf{C}}_W^{i,j}(\boldsymbol{\theta})$

5.2 Computation of the multi-level matrix blocks of $\tilde{\mathbf{C}}_W$

Given an octree multi-level tree domain decomposition in $\mathbb{R}^3$, as shown in [8, 9], the authors described how to apply a Kernel Independent Fast Multipole Method (KIFMM) by [32] to compute all the blocks $\tilde{\mathbf{C}}_W^{i,i}(\boldsymbol{\theta}) \in \mathbb{R}^{\tilde{N} \times \tilde{N}}$ for $i = -1, \ldots, t$ in $\mathcal{O}(\tilde{N}(t+1)^2)$ computational steps to a fixed accuracy $\varepsilon_{FMM} > 0$.

For the random projection tree it is not possible to determine a-priori the sparsity of the sparse blocks $\tilde{\mathbf{C}}_W^{i,j}(\boldsymbol{\theta})$. But given a value for τ by running Algorithm 7 on every cell $B_k^i \in \mathcal{B}^i$, at level

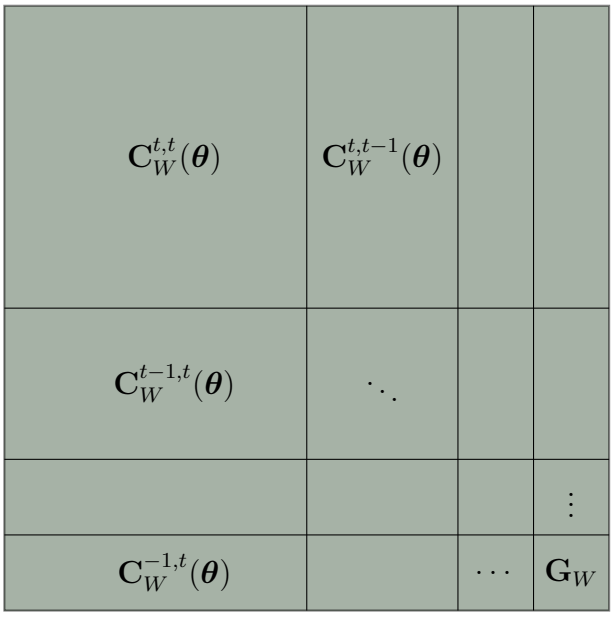

Figure 4: Multi-Level covariance matrix where $\mathbf{G}_W := \mathbf{C}_W^{-1,-1}(\boldsymbol{\theta})$.

i, with the Targetdepth corresponding for level j it is possible to determine the computational cost of constructing the sparse blocks $\tilde{\mathbf{C}}_W^{i,j}(\boldsymbol{\theta})$. Suppose that maximum number of cells $B_k^j \in \mathcal{B}^j$ given by Algorithm 7 is bounded by some $\gamma^{i,j} \in \mathbb{N}_+$.

Assumption 1. *Let $\mathbf{A}(\boldsymbol{\theta}) \in \mathbb{R}^{\tilde{M}\times\tilde{N}}$ be a kernel matrix formed from $\tilde{N}$ source observation nodes and $\tilde{M}$ target nodes in the space $\mathbb{R}^d$. Suppose that there exists a fast summation method that computes the matrix-vector products $\mathbf{A}(\boldsymbol{\theta})\mathbf{x}$ with $\varepsilon_{FMM} > 0$ accuracy in $\mathcal{O}((\tilde{N}+\tilde{M})^\alpha)$ computations, for some $\alpha \geq 1$ and any $\mathbf{x} \in \mathbb{R}^d$.*

For $\mathbb{R}^3$ problems the Kernel Independent Fast Multipole Method (KIFMM) is used. This method is flexible and efficient for computing the matrix vector products $\mathbf{C}(\boldsymbol{\theta})\mathbf{x}$ for a large class of kernel functions, including the Matérn covariance function. Experimental results show a computational cost of about $\mathcal{O}(\tilde{N}+\tilde{M})$, $\alpha \approx 1$ with good accuracy (ε_{FMM} between 10^{-6} to 10^{-8}) with a slight degrade in the accuracy with increased source nodes.

Now, for problems where $d > 3$, an Approximate Skeletonization Kernel Independent Treecode (ASKIT) method [22] is used instead to compute the matrix-vector products. The numerical results for this approach are mixed. For $\nu \to \infty$, i.e. Gaussian kernel, ASKIT exhibits a computational cost of about $\mathcal{O}(\tilde{N}+\tilde{M})$, with high accuracy $\varepsilon_{FMM} \approx 10^{-8}$. However, if $N > 32,000$ the observed computational cost is only slightly better than the direct method, but still degrades as $\mathcal{O}(\tilde{N}\tilde{M})$. Now, for the Matérn kernel with $\nu < \infty$, the observed computational cost is about $\mathcal{O}(\tilde{N}\tilde{M})$.

Theorem 4. *The cost of computing each block $\tilde{\mathbf{C}}_W^{i,j}(\boldsymbol{\theta})$ for $i,j = 1,\ldots,t$ by using a fast summation method with $1 \leq \alpha \leq 2$ is bounded by*

$$\mathcal{O}(\gamma_{i,j}\tilde{p}2^i(2^{t-j+1}\tilde{p} + 2^{t-i+1}\tilde{p})^\alpha + 2\tilde{p}2^t)$$

Proof. Let us look at the cost of computing all the interactions between any two cells $B_k^i \in \mathcal{B}^i$ and $B_l^j \in \mathcal{B}^j$. Without loss of generality assume that $i \leq j$. For the cell B_k^l there is at most

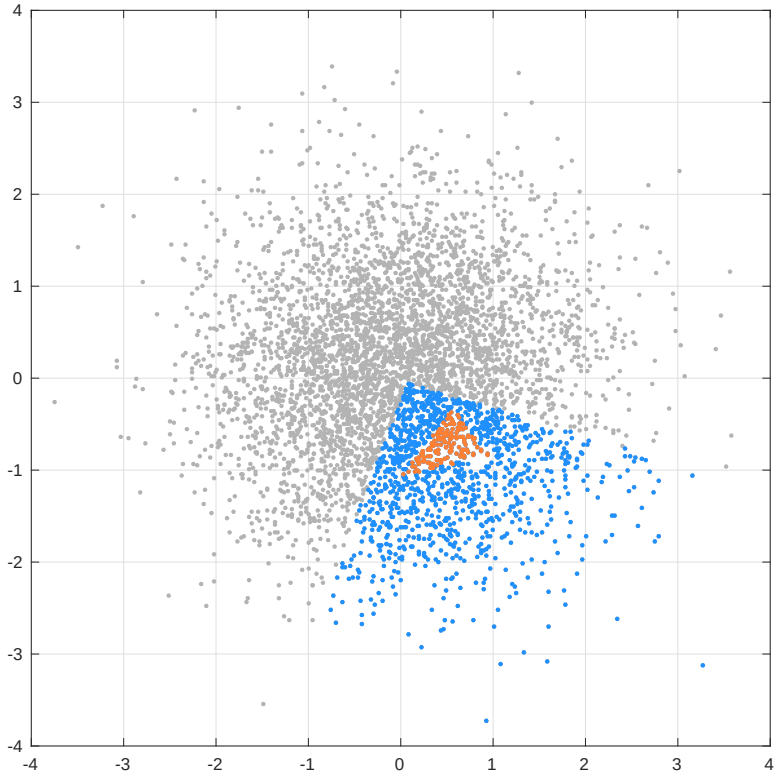

Figure 5: Neighborhood identification from source cell on a random projection tree decomposition of observation locations in $\mathbb{R}^2$. The orange nodes are contained in the source cell. The blue nodes are are contained in the local neighborhood cells. The grey dots are all the observations that are not part of the source or local neighborhood cells.

$\tilde{p}$ multi-level vectors and from Lemma 2 $2^{t-i+1}\tilde{p}$ non zero entries. Similarly for B_l^j. All the interactions $(\boldsymbol{\psi}_k^i)^{\mathrm{T}}\mathbf{C}(\boldsymbol{\theta})\boldsymbol{\psi}_l^j$ now have to be computed, where $\boldsymbol{\psi}_k^i \in B_k^i$ and $\boldsymbol{\psi}_l^j \in B_l^j$.

The term $\mathbf{C}(\boldsymbol{\theta})\boldsymbol{\psi}_l^j$ is computed by using an FMM with $2^{t-j+1}\tilde{p}$ sources and $2^{t-i+1}\tilde{p}$ targets at a cost of $\mathcal{O}((2^{t-j+1}\tilde{p} + 2^{t-i+1}\tilde{p})^{\alpha})$. Since there is at most $\tilde{p}$ multi-level vectors in B_k^i and B_l^j then the cost for computing all the interactions $(\boldsymbol{\psi}_k^i)^{\mathrm{T}}\mathbf{C}(\boldsymbol{\theta})\boldsymbol{\psi}_l^j$ is $\mathcal{O}(\tilde{p}(2^{t-j+1}\tilde{p} + 2^{t-i+1}\tilde{p})^{\alpha} + 2^{t-i+1}\tilde{p})$.

Now, at any level i there is at most 2^i cells, thus the result follows. □

6 Multi-level estimator and predictor

The multi-level random projection tree can be exploited in such a way to significantly reduce the computational burden and to further increase the numerical stability of the estimation and prediction steps. This is a direct extension of the multi-level estimator and predictor formulated in [8] for higher dimensions.

6.1 Estimator

The multi-level likelihood function, $l_W(\theta)$ (see equation (12)), has the clear advantage of being decoupled from the vector $\boldsymbol{\beta}$. Furthermore, the multi-level covariance matrix $\mathbf{C}_W(\boldsymbol{\theta})$ will be more numerically stable than $\mathbf{C}(\boldsymbol{\theta})$ thus making it easier to invert and to compute the determinant.

However, it is not necessary to perform the MLE estimation on the full covariance matrix $\mathbf{C}_W(\boldsymbol{\theta})$, instead construct a series of multi-level likelihood functions $\tilde{\ell}_W^n(\boldsymbol{\theta})$, $n = -1, \dots t$, by applying the partial transform $[\mathbf{W}_t^{\mathrm{T}}, \dots, \mathbf{W}_n^{\mathrm{T}}]$ to the data $\mathbf{Z}$ where

$$\tilde{\ell}_W^n(\boldsymbol{\theta}) = -\frac{\tilde{N}}{2}\log(2\pi) - \frac{1}{2}\log\det\{\tilde{\mathbf{C}}_W^n(\boldsymbol{\theta})\} - \frac{1}{2}(\mathbf{Z}_W^n)^{\mathrm{T}}\tilde{\mathbf{C}}_W^n(\boldsymbol{\theta})^{-1}\mathbf{Z}_W^n, \tag{14}$$

where $\mathbf{Z}_W^n := [\mathbf{W}_t^{\mathrm{T}}, \dots, \mathbf{W}_i^{\mathrm{T}}]^{\mathrm{T}}\mathbf{Z}$, $\tilde{N}$ is the length of $\mathbf{Z}_W^n$, $\tilde{\mathbf{C}}_W^n(\boldsymbol{\theta})$ is the $\tilde{N}\times\tilde{N}$ upper-left sub-matrix of $\tilde{\mathbf{C}}_W(\boldsymbol{\theta})$ and $\mathbf{C}_W^n(\boldsymbol{\theta})$ is the $\tilde{N}\times\tilde{N}$ upper-left sub-matrix of $\mathbf{C}_W(\boldsymbol{\theta})$.

A consequence of this approach is that the matrices $\mathbf{C}_W^n(\boldsymbol{\theta})$, $n = -1, \dots, t$ will be increasingly more numerically stable, thus easier to solve computationally, as shown in the following theorem.

Theorem 5. *Let $\kappa(A) \to \mathbb{R}$ be the condition number of the matrix $A \in \mathbb{R}^{N\times N}$ then*

$$\kappa(\mathbf{C}_{\mathbf{W}}^t(\boldsymbol{\theta})) \le \kappa(\mathbf{C}_{\mathbf{W}}^{t-1}(\boldsymbol{\theta})) \le \cdots \le \kappa(\mathbf{C}_{\mathbf{W}}(\boldsymbol{\theta})) \le \kappa(\mathbf{C}(\boldsymbol{\theta})).$$

Proof. This proof is a simple extension of the proof in Theorem 3. □

Remark 11. *If $\mathbf{C}(\boldsymbol{\theta})$ is symmetric positive definite then for $n = -1, \dots, t$ the matrices $\mathbf{C}_W^n(\boldsymbol{\theta})$ are symmetric positive definite. The proof is immediate.*

Remark 12. *If the matrix $\tilde{\mathbf{C}}_W^n(\boldsymbol{\theta})$ is close to $\mathbf{C}_W^n(\boldsymbol{\theta})$, for $n = 1, \dots, d$, in some matrix norm sense, the condition number of $\tilde{\mathbf{C}}_W^n(\boldsymbol{\theta})$ will be close to $\mathbf{C}_W^n(\boldsymbol{\theta})$. In section 7, for a class of covariance functions, it can be shown that for sufficiently large τ and/or $w, a \in \mathbb{N}_+$ [with the index set $\Lambda(w)^{m,g}$ or $\tilde{\Lambda}(w)^{m,g}$] $\tilde{\mathbf{C}}_W^n(\boldsymbol{\theta})$ will be close to $\mathbf{C}_W^n(\boldsymbol{\theta})$. Thus $\tilde{\mathbf{C}}_W^n(\boldsymbol{\theta})$ will be symmetric positive definite.*

6.2 Predictor

An alternative formulation for obtaining the estimate $\hat{Z}(\mathbf{x}_0)$ is by solving the following problem

$$\begin{pmatrix} \mathbf{C}(\boldsymbol{\theta}) & \mathbf{M} \\ \mathbf{M}^{\mathrm{T}} & \mathbf{0} \end{pmatrix}\begin{pmatrix} \hat{\boldsymbol{\gamma}} \\ \hat{\boldsymbol{\beta}} \end{pmatrix} = \begin{pmatrix} \mathbf{Z} \\ \mathbf{0} \end{pmatrix}. \tag{15}$$

In [25], the authors show that the solution of this problem leads to equation (4) and $\hat{\boldsymbol{\gamma}}(\boldsymbol{\theta}) = \mathbf{C}^{-1}(\boldsymbol{\theta})(\mathbf{Z} - \mathbf{M}\hat{\boldsymbol{\beta}}(\boldsymbol{\theta}))$. The best unbiased predictor can be evaluated as

$$\hat{Z}(\mathbf{x}_0) = \mathbf{k}(\mathbf{x}_0)^{\mathrm{T}}\hat{\boldsymbol{\beta}}(\boldsymbol{\theta}) + \mathbf{c}(\boldsymbol{\theta})^{\mathrm{T}}\hat{\boldsymbol{\gamma}}(\boldsymbol{\theta}) \tag{16}$$

and the Mean Squared Error (MSE) at the target point $\mathbf{x}_0$ is given by

$$1 + \tilde{\mathbf{u}}^{\mathrm{T}}(\mathbf{M}^{\mathrm{T}}\mathbf{C}(\boldsymbol{\theta})^{-1}\mathbf{M})^{-1}\tilde{\mathbf{u}} - \mathbf{c}(\boldsymbol{\theta})^{\mathrm{T}}\mathbf{C}^{-1}(\boldsymbol{\theta})\mathbf{c}(\boldsymbol{\theta})$$

where $\tilde{\mathbf{u}}^{\mathrm{T}} := (\mathbf{M}\mathbf{C}^{-1}(\boldsymbol{\theta})\mathbf{c}(\boldsymbol{\theta}) - \mathbf{k}(\mathbf{x}_0))$.

From (15) it is observed that $\mathbf{M}^{\mathrm{T}}\hat{\boldsymbol{\gamma}}(\boldsymbol{\theta}) = \mathbf{0}$. This implies that $\hat{\boldsymbol{\gamma}} \in \mathbb{R}^n \backslash \mathcal{P}^p(\mathbb{S})$ and can be uniquely rewritten as $\hat{\boldsymbol{\gamma}} = \mathbf{W}^{\mathrm{T}}\boldsymbol{\gamma}_W$ for some $\boldsymbol{\gamma}_W \in \mathbb{R}^{N-p}$. Now, rewrite $\mathbf{C}(\boldsymbol{\theta})\hat{\boldsymbol{\gamma}} + \mathbf{M}\hat{\boldsymbol{\beta}} = \mathbf{Z}$ as

$$\mathbf{C}(\boldsymbol{\theta})\mathbf{W}^{\mathrm{T}}\hat{\boldsymbol{\gamma}}_W + \mathbf{M}\hat{\boldsymbol{\beta}} = \mathbf{Z}. \tag{17}$$

Now apply the matrix $\mathbf{W}$ to equation (17) and obtain $\mathbf{W}\{\mathbf{C}(\boldsymbol{\theta})\mathbf{W}^{\mathrm{T}}\hat{\boldsymbol{\gamma}}_W + \mathbf{M}\hat{\boldsymbol{\beta}}\} = \mathbf{W}\mathbf{Z}$. Since $\mathbf{WM} = \mathbf{0}$ then

$$\mathbf{C}_W(\boldsymbol{\theta})\hat{\boldsymbol{\gamma}}_W = \mathbf{Z}_W.$$

The preconditioner $\mathbf{P}_W$ is formed from the entries of the matrix $\mathbf{C}_W$. For $i = -1, \ldots, t$ and each cell $B_k^i \in \mathcal{B}^i$ form a matrix $\mathbf{P}^{i,k}$ of all the interactions $(\boldsymbol{\psi}_{\tilde{k}}^i)^{\mathrm{T}}\mathbf{C}(\boldsymbol{\theta})\boldsymbol{\psi}_{\tilde{k}}^i$, where $\boldsymbol{\psi}_{\tilde{k}}^i \in B_k^i$. Thus,

$$\mathbf{P}^{i,k} := [\boldsymbol{\psi}_1^i, \ldots, \boldsymbol{\psi}_s^i]^{\mathrm{T}}\mathbf{C}_W[\boldsymbol{\psi}_1^i, \ldots, \boldsymbol{\psi}_s^i]$$

where $\boldsymbol{\psi}_1^i, \ldots, \boldsymbol{\psi}_s^i \in B_k^i$ and s is the total number of ML basis in the cell B_k^i. The preconditioner is formed as a block diagonal matrix of all the matrices $\mathbf{P}^{i,k}$ for every cell $B_k^i \in \mathcal{B}$, i.e.

$$\mathbf{P}_W := diag(\mathbf{P}^{-1,0}, \mathbf{P}^{0,0}, \mathbf{P}^{1,0}, \mathbf{P}^{1,1}, \ldots).$$

Applying the preconditioner $\mathbf{P}_W^{-1}(\boldsymbol{\theta})$ the following system of equations

$$\bar{\mathbf{C}}_W(\boldsymbol{\theta})\boldsymbol{\gamma}_W(\boldsymbol{\theta}) = \bar{\mathbf{Z}}_W$$

are obtained, where $\bar{\mathbf{C}}_W(\boldsymbol{\theta}) := \mathbf{P}_W^{-1}(\boldsymbol{\theta})\mathbf{C}_W(\boldsymbol{\theta})$ and $\bar{\mathbf{Z}}_W := \mathbf{P}_W^{-1}(\boldsymbol{\theta})\mathbf{Z}_W$. Note that in some cases $\mathbf{C}_W(\boldsymbol{\theta})$ will have very small condition numbers. For this case we can set $\mathbf{P}_W := I$, i.e. no preconditioner.

Theorem 6. *If the covariance function $\phi : \Gamma_d \times \Gamma_d \to \mathbb{R}$ is positive definite, then the matrix $\mathbf{P}_W(\boldsymbol{\theta})$ is always symmetric positive definite.*

Proof. Immediate. □

6.3 Estimator: Computation of $\log\det\{\tilde{\mathbf{C}}_W^n\}$ and $(\mathbf{Z}_W^n)^{\mathrm{T}}(\tilde{\mathbf{C}}_W^n)^{-1}\mathbf{Z}_W^n$

An approach to computing the determinant of $\tilde{\mathbf{C}}_W^n(\boldsymbol{\theta})$ is to apply a sparse Cholesky factorization technique such that $\mathbf{G}\mathbf{G}^{\mathrm{T}} = \tilde{\mathbf{C}}_W^n(\boldsymbol{\theta})$, where $\mathbf{G}$ is a lower triangular matrix. Since the eigenvalues of $\mathbf{G}$ are located on the diagonal then $\log\det\{\tilde{\mathbf{C}}_W^n(\boldsymbol{\theta})\} = 2\sum_{i=1}^{\tilde{N}} \log \mathbf{G}_{ii}$.

To reduce the fill-in of the factorization matrix $\mathbf{G}$ apply the matrix reordering technique in Suite Sparse 4.2.1 package ([11, 13–16]) with the Nested Dissection (NESDIS) function package. The sparse Cholesky factorization is performed with the *lchol* command from Suite Sparse 4.2.1 package.

Although in practice the combined NESDIS and sparse Cholesky factorization is highly efficient, as shown by our numerical results, a worse case complexity bound can be obtained only for d = 2 or d = 3 dimensions (see [8]).

Two choices for the computation of $(\mathbf{Z}_W^n)^{\mathrm{T}}\tilde{\mathbf{C}}_W^n(\boldsymbol{\theta})^{-1}\tilde{\mathbf{Z}}_W^n$ are open to us: i) a Cholesky factorization of $\tilde{\mathbf{C}}_W^n(\boldsymbol{\theta})$, or ii) a Preconditioned Conjugate Gradient (PCG). The PCG choice requires significantly less memory and allows more control of the error. However, the sparse Cholesky factorization of $\tilde{\mathbf{C}}_W^n(\boldsymbol{\theta})$ has already been used to compute the determinant. Thus we can use the same factors to compute $(\tilde{\mathbf{Z}}_W^n)^{\mathrm{T}}\tilde{\mathbf{C}}_W^n(\boldsymbol{\theta})^{-1}\tilde{\mathbf{Z}}_W^n$.

6.4 Predictor: Computation of $\bar{\mathbf{C}}_W\hat{\boldsymbol{\gamma}}_W = \bar{\mathbf{Z}}_W$

For the predictor stage a different approach is used. Instead of inverting the sparse matrix $\tilde{C}_W(\boldsymbol{\theta})$ a Preconditioned Conjugate Gradient (PCG) method is employed to compute $\hat{\boldsymbol{\gamma}}_W = \mathbf{C}_W(\boldsymbol{\theta})^{-1}\mathbf{Z}_W$.

Recall that $\mathbf{C}_W = \mathbf{W}\mathbf{C}(\boldsymbol{\theta})\mathbf{W}^{\mathrm{T}}$, $\hat{\boldsymbol{\gamma}}_W = \mathbf{W}\hat{\boldsymbol{\gamma}}$ and $\mathbf{Z}_W = \mathbf{W}\mathbf{Z}$. Thus the matrix vector products $\mathbf{C}_W(\boldsymbol{\theta})\boldsymbol{\gamma}_W^n$ in the PCG iteration are computed within three steps:

$$\boldsymbol{\gamma}_W^n \xrightarrow[(1)]{\mathbf{W}^{\mathrm{T}}\boldsymbol{\gamma}_W^n} \mathbf{a}_n \xrightarrow[(2)]{\mathbf{C}(\boldsymbol{\theta})\mathbf{a}_n} \mathbf{b}_n \xrightarrow[(3)]{\mathbf{W}\mathbf{b}_n} \mathbf{C}_W(\boldsymbol{\theta})\boldsymbol{\gamma}_W^n$$

where $\boldsymbol{\gamma}_W^0$ is the initial guess and $\boldsymbol{\gamma}_W^n$ is the n^{th} iteration of the PCG.

(1) Transformation from multi-level representation to single level. This is done in at most $\mathcal{O}(Nt)$ steps.

(2) Perform matrix vector product using a fast summation method. For $d = 2, 3$ a KIFMM is used to compute the matrix vector products. Alternatively, for $d > 3$ an ASKIT method is used to compute the matrix-vector products.

(3) Convert back to multi-level representation.

The matrix-vector products $\mathbf{C}_W(\boldsymbol{\theta})\boldsymbol{\gamma}_W^n$, where $\boldsymbol{\gamma}_W^n \in \mathbb{R}^{N-p}$, are computed in $\mathcal{O}(N^\alpha + 2Nt)$ computational steps to a fixed accuracy $\varepsilon_{FMM} > 0$. Note that $\alpha \geq 1$ is dependent on the efficiency of the fast summation method. The total computational cost is $\mathcal{O}(kN^\alpha + 2Nt)$, where k is the number of iterations needed to solve $\bar{\mathbf{C}}_W(\boldsymbol{\theta})\bar{\boldsymbol{\gamma}}_W(\boldsymbol{\theta}) = \bar{\mathbf{Z}}_W$ to a predetermined accuracy $\varepsilon_{PCG} > 0$.

It is important to point out that the introduction of a preconditioner can degrade the performance of the PCG, in particular, if the preconditioner is ill-conditioned. The accuracy of the PCG method ε_{PCG} has to be set such that the accuracy of the *unpreconditioned* system $\mathbf{C}_W(\boldsymbol{\theta})\boldsymbol{\gamma}_W(\boldsymbol{\theta}) = \mathbf{Z}_W$ is below a user given tolerance $\varepsilon > 0$.

Now compute $\hat{\boldsymbol{\gamma}} = \mathbf{W}^{\mathrm{T}}\hat{\boldsymbol{\gamma}}_W$ and $\hat{\boldsymbol{\beta}} = (\mathbf{M}^{\mathrm{T}}\mathbf{M})^{-1}\mathbf{M}^{\mathrm{T}}(\mathbf{Z} - \mathbf{C}(\boldsymbol{\theta})\hat{\boldsymbol{\gamma}})$ in at most $\mathcal{O}(N^\alpha + Np + p^3)$ computational steps. The matrix vector product $\mathbf{c}(\boldsymbol{\theta})^{\mathrm{T}}\hat{\boldsymbol{\gamma}}(\boldsymbol{\theta})$ is computed in $\mathcal{O}(N)$ steps. Finally, the total cost for computing the estimate $\hat{\mathbf{Z}}(\mathbf{x}_0)$ from (16) is $\mathcal{O}(p^3 + (k+1)N^\alpha + 2Nt)$.

7 Error analysis

Consider the full solution $\mathbf{x}_W = \mathbf{C}_W^n(\boldsymbol{\theta})^{-1}\mathbf{Z}_W^n$ and sparse solution $\tilde{\mathbf{x}}_W = \tilde{\mathbf{C}}_W^n(\boldsymbol{\theta})^{-1}\mathbf{Z}_W^n$ for $n = -1, \ldots, t$, then the error can be bounded as

$$\|\mathbf{x}_W - \tilde{\mathbf{x}}_W\|_{l^2} \leq \|\mathbf{C}_W^{-1}(\boldsymbol{\theta}) - \tilde{\mathbf{C}}_W^{-1}(\boldsymbol{\theta})\|_2 \|\mathbf{Z}_W\|_{l^2}. \tag{18}$$

This section is divided into three parts: i) decay of the multi-level matrix $\mathbf{C}_W(\boldsymbol{\theta})$ coefficients ii) analyticity of the Matérn covariance function and finally iii) a-posteriori error estimates of $\|\mathbf{x}_W - \tilde{\mathbf{x}}_W\|_{l^2}$.

7.1 Multi-level covariance matrix decay

The decay of the coefficients of the matrix $\mathbf{C}_W(\boldsymbol{\theta})$ will depend directly on the choice of the multivariate index set $\Lambda^{m,g}(w)$ and the analytic regularity of the covariance function. In this section two lemmas are derived that provide bounds on the magnitude of the coefficients of $\mathbf{C}_W(\boldsymbol{\theta})$.

For the case that $\nu \to \infty$ the Matérn covariance function converges to a Gaussian and is analytic on $\Gamma_d \times \Gamma_d$. For this case error bounds that depend on the domain $\Gamma_d \times \Gamma_d$ are derived

in Lemma 3 with respect to the subspace $\mathcal{P}^{\tilde{p}}(\mathbb{S})$ generated by the Smolyak index set $\Lambda^{m,g}(w+a)$, where $w, a \in \mathbb{N}_+$.

In general, the Matérn covariance function will not be regular at the origin. If that is the case, care has to be taken that the sparse grid domain does not cross the origin. Error bounds that depend on the extended Smolyak index set $\tilde{\Lambda}^{m,g}(w+a)$, where $w, a \in \mathbb{N}_+$ are derived in Lemma 4.

Lemma 3. *Suppose that $\phi(\mathbf{x}, \mathbf{y}; \boldsymbol{\theta}) \in C^0(\Gamma^d \times \Gamma^d; \mathbb{R})$ has an analytic extension on $\mathcal{E}_{\hat{\sigma},\ldots,\hat{\sigma}} \times \mathcal{E}_{\hat{\sigma},\ldots,\hat{\sigma}}$ and is bounded by $\tilde{M}(\phi)$. Let $\mathcal{P}^{\tilde{p}}(\mathbb{S})^{\perp}$ be the subspace in $\mathbb{R}^N$ generated by the Smolyak index set $\Lambda^{m,g}(w+a)$ for some $w, a \in \mathbb{N}_+$. Let $\mathcal{S}^{m,g}_{w+a,d}$ be a sparse grid generated from the index set $\Lambda^{m,g}(w+a)$ with $\eta(w+a)$ collocation nodes. For $i, j = 0, \ldots, t$ consider any multi-level vector $\boldsymbol{\psi}^i_m \in \mathcal{P}^{\tilde{p}}(\mathbb{S})^{\perp}$, with n_m non-zero entries, from the cell $B^i_m \in \mathcal{B}^i$ and any multi-level vector $\boldsymbol{\psi}^j_q \in \mathcal{P}^{\tilde{p}}(\mathbb{S})^{\perp}$, with n_q non-zero entries, from the cell $B^j_q \in \mathcal{B}^j$, then:*

(a) if $w + a > d/\log(2)$

$$\left| \sum_{r=1}^{N} \sum_{h=1}^{N} \phi(\mathbf{x}_r, \mathbf{y}_h; \boldsymbol{\theta}) \boldsymbol{\psi}^i_m[h] \boldsymbol{\psi}^j_q[r] \right| \leq \sqrt{n_m n_q} \mathcal{M}(\sigma, \delta^*(\sigma), d, \tilde{M}(\phi)) \eta^{2\mu_3(\sigma,\delta^*(\sigma),d)} \exp\left(-\frac{2d\sigma}{2^{1/d}} \eta^{\mu_2(d)} \right),$$

where

$$\mathcal{M}(\sigma, \delta^*(\sigma), d, \tilde{M}_{\epsilon_1}, \tilde{M}(\phi)) := \frac{8\mathcal{Q}(\sigma, \delta^*(\sigma), d, \tilde{M}(\phi)) \tilde{C}(\sigma) a(\delta, \sigma)}{\exp(\sigma \delta^*(\sigma) \tilde{C}_2(\sigma)) e \delta \sigma} \frac{\max\{1, \tilde{C}_1(\sigma, \delta^*(\sigma), d, \tilde{M}_{\epsilon_1})\}}{|1 - \tilde{C}_1(\sigma, \delta^*(\sigma), d, \tilde{M}_{\epsilon_1})|},$$

and

$$\tilde{M}_{\epsilon_1}(\sigma, \delta^*(\sigma), d, \tilde{M}(\phi)) := \tilde{\mathcal{Q}}(\sigma, \delta^*(\sigma), d, \tilde{M}(\phi)) \eta^{\mu_3(\sigma,\delta^*(\sigma),d)} \exp\left(-\frac{d\sigma}{2^{1/d}} \eta^{\mu_2(d)} \right).$$

(b) if $w + a \leq d/\log(2)$

$$\left| \sum_{r=1}^{N} \sum_{h=1}^{N} \phi(\mathbf{x}_r, \mathbf{y}_h; \boldsymbol{\theta}) \boldsymbol{\psi}^i_m[h] \boldsymbol{\psi}^j_q[r] \right| \leq \sqrt{n_m n_q} \frac{8C(\sigma) a(\delta, \sigma) C_1(\sigma, \delta^*(\sigma), d, \tilde{M}(\phi))}{e \delta \sigma} \eta^{-2\mu_1}$$
$$* \frac{\max\{1, C_1(\sigma, \delta^*(\sigma), d, \tilde{P}_{\epsilon_1})\}^d}{|1 - C_1(\sigma, \delta^*(\sigma), d, \tilde{P}_{\epsilon_1})|} \frac{\max\{1, C_1(\sigma, \delta^*(\sigma), d, \tilde{M}(\phi)\}^d}{|1 - C_1(\sigma, \delta^*(\sigma), d, \tilde{M}(\phi))|}.$$

where

$$\tilde{P}_{\epsilon_1}(\sigma, \delta^*(\sigma), d, \tilde{M}(\phi)) := \frac{2C_1(\sigma, \delta^*(\sigma), \tilde{M}(\phi))}{|1 - C_1(\sigma, \delta^*(\sigma), \tilde{M}(\phi))|} \max\{1, C_1(\sigma, \delta^*(\sigma), \tilde{M}(\phi))\}^d \eta^{-\mu_1}.$$

Proof. Let $\epsilon_1(\mathbf{x}, \mathbf{y}; \boldsymbol{\theta}) := (I_d \otimes I_d - \mathcal{S}^{m,g}_{w+a,d} \otimes I_d)\phi(\mathbf{x}, \mathbf{y}; \boldsymbol{\theta})$, where $I_d : C^d(\Gamma_d; \mathbb{R}) \to C^d(\Gamma_d; \mathbb{R})$ is the identity operator, and $\epsilon_2(\mathbf{x}, \mathbf{y}; \boldsymbol{\theta}) := (I_d \otimes I_d - I_d \otimes \mathcal{S}^{m,g}_{w+a,d})\epsilon_1(\mathbf{x}, \mathbf{y}; \boldsymbol{\theta})$. Consider any multi-level vector $\boldsymbol{\psi}^i_m \in \mathcal{P}^{\tilde{p}}(\mathbb{S})^{\perp}$ in the cell B^m and any vector $\boldsymbol{\psi}^j_q \in \mathcal{P}^{\tilde{p}}(\mathbb{S})^{\perp}$ from the cell B^q, where

$m, q = 0, \ldots, t$, then

$$\begin{aligned}\sum_{k=1}^{N}\sum_{l=1}^{N}\phi(\mathbf{x}_k,\mathbf{y}_l;\boldsymbol{\theta})\boldsymbol{\psi}_m^i[k]\boldsymbol{\psi}_q^j[l] &= \sum_{k=1}^{N}\boldsymbol{\psi}_m^i[k]\left[\sum_{l=1}^{N}\phi(\mathbf{x}_k,\mathbf{y}_l;\boldsymbol{\theta})\boldsymbol{\psi}_q^j[l]\right]\\ &= \sum_{k=1}^{N}\boldsymbol{\psi}_m^i[k]\left[\sum_{l=1}^{N}\epsilon_1(\mathbf{x}_k,\mathbf{y}_l;\boldsymbol{\theta}) - (I_d\otimes\mathcal{S}_{w+a,d}^{m,g})\phi(\mathbf{x}_k,\mathbf{y}_l;\boldsymbol{\theta})\boldsymbol{\psi}_q^j[l]\right]\\ &= \sum_{k=1}^{N}\boldsymbol{\psi}_m^i[k]\left[\sum_{l=1}^{N}\epsilon_1(\mathbf{x}_k,\mathbf{y}_l;\boldsymbol{\theta})\boldsymbol{\psi}_q^j[l]\right].\end{aligned}$$

The last equality follows from the fact that from Proposition 1 $(I_d\otimes\mathcal{S}_{w+a,d}^{m,g})\phi(\mathbf{x},\mathbf{y};\boldsymbol{\theta})\in I_d\otimes C^0(\Gamma_d;\mathbb{R})\otimes\mathbb{P}_{\Lambda^{m,g}(w)}$. Since $\boldsymbol{\psi}_j^q\in\mathcal{P}^{\tilde{p}}(\mathbb{S})^{\perp}$ then $\sum_{l=1}^{N}(I_d\otimes\mathcal{S}_{w+a,d}^{m,g})\phi(\mathbf{x}_k,\mathbf{y}_l;\boldsymbol{\theta})\boldsymbol{\psi}_q^j[l]=0$. Now, rearranging the summations:

$$\begin{aligned}\sum_{k=1}^{N}\sum_{l=1}^{N}\phi(\mathbf{x}_k,\mathbf{y}_l;\boldsymbol{\theta})\boldsymbol{\psi}_m^i[k]\boldsymbol{\psi}_q^j[l] &= \sum_{l=1}^{N}\boldsymbol{\psi}_q^j[k]\left[\sum_{k=1}^{N}\epsilon_1(\mathbf{x}_k,\mathbf{y}_l;\boldsymbol{\theta})\boldsymbol{\psi}_m^i[l]\right]\\ &= \sum_{l=1}^{N}\boldsymbol{\psi}_q^j[k]\left[\sum_{k=1}^{N}\epsilon_2(\mathbf{x}_k,\mathbf{y}_l;\boldsymbol{\theta}) + (\mathcal{S}_{w+a,d}^{m,g}\otimes I_d)\epsilon_1(\mathbf{x}_k,\mathbf{y}_l;\boldsymbol{\theta})\boldsymbol{\psi}_m^i[l]\right]\\ &= \sum_{l=1}^{N}\boldsymbol{\psi}_q^j[k]\left[\sum_{k=1}^{N}\epsilon_2(\mathbf{x}_k,\mathbf{y}_l;\boldsymbol{\theta})\boldsymbol{\psi}_m^i[l]\right]\\ &= \sum_{k=1}^{N}\sum_{l=1}^{N}\epsilon_2(\mathbf{x}_k,\mathbf{y}_l;\boldsymbol{\theta})\boldsymbol{\psi}_m^i[k]\boldsymbol{\psi}_q^j[l].\end{aligned}$$

Now,

$$\begin{aligned}\left|\sum_{r=1}^{N}\sum_{h=1}^{N}\phi(\mathbf{x}_r,\mathbf{y}_h;\boldsymbol{\theta})\boldsymbol{\psi}_m^i[h]\boldsymbol{\psi}_q^j[r]\right| &\leq \|\epsilon_2(\mathbf{x},\mathbf{y};\boldsymbol{\theta})\|_{L^\infty(\Gamma_d\times\Gamma_d)}\sum_{r=1}^{N}\sum_{h=1}^{N}|\boldsymbol{\psi}_m^i[h]\boldsymbol{\psi}_q^j[r]|\\ &\leq \|\epsilon_2(\mathbf{x},\mathbf{y};\boldsymbol{\theta})\|_{L^\infty(\Gamma_d\times\Gamma_d)}\sqrt{n_m}\sqrt{n_q}\|\boldsymbol{\psi}_m^i\|_{l^2}\|\boldsymbol{\psi}_q^j\|_{l^2}\\ &= \|\epsilon_2(\mathbf{x},\mathbf{y};\boldsymbol{\theta})\|_{L^\infty(\Gamma_d\times\Gamma_d)}\sqrt{n_m}\sqrt{n_q}.\end{aligned}$$

The last equality is due to the fact that $\boldsymbol{\psi}_m^i$ and $\boldsymbol{\psi}_q^j$ are orthonormal. The next step is to bound the error $\|\epsilon_2(\mathbf{x},\mathbf{y};\boldsymbol{\theta})\|_{L^\infty(\Gamma_d\times\Gamma_d)}$.

Now, extend $\mathbf{y}\to\mathbf{z}\in\mathcal{E}_{\hat{\sigma},\ldots,\hat{\sigma}}$. For this case $\epsilon_1(\cdot,\mathbf{z};\boldsymbol{\theta})$ becomes a complex number. If $w > d/\log(2)$ then from Theorem 2 and equation (11) for all $\mathbf{z}\in\mathcal{E}_{\hat{\sigma},\ldots,\hat{\sigma}}$

$$\begin{aligned}\|\epsilon_1(\cdot,\mathbf{z};\boldsymbol{\theta})\|_{L^\infty(\Gamma^d)} &= \|(I_d-\mathcal{S}_{w+a,d}^{m,g})\operatorname{Re}\phi(\cdot,\mathbf{z}) + i(I_d-\mathcal{S}_{w+a,d}^{m,g})\operatorname{Im}\phi(\cdot,\mathbf{z})\|_{L^\infty(\Gamma^d)}\\ &\leq \tilde{M}_{\epsilon_1} := 2\mathcal{Q}(\sigma,\delta^*(\sigma),d,\tilde{M}(\phi)\eta^{\mu_3(\sigma,\delta^*(\sigma),d)}\exp\left(-\frac{d\sigma}{2^{1/d}}\eta^{\mu_2(d)}\right),\end{aligned}$$

where $\tilde{M}(\phi) = \max_{\mathbf{w},\mathbf{z}\in\mathcal{E}_{\hat{\sigma},\ldots,\hat{\sigma}}\times\mathcal{E}_{\hat{\sigma},\ldots,\hat{\sigma}}} |\phi(\mathbf{w},\mathbf{z})|$. From Theorem 7 and equation (10) it follows that

$$\begin{aligned}\|\epsilon_2(\mathbf{x},\mathbf{y};\boldsymbol{\theta})\|_{L^\infty(\Gamma^d\times\Gamma^d)} &= \|(I_d\otimes I_d - I_d\otimes\mathcal{S}^{m,g}_{w+a,d})\epsilon_1(\mathbf{x},\mathbf{y};\boldsymbol{\theta})\|_{L^\infty(\Gamma^d\times\Gamma^d)}\\ &= \|(I_d\otimes(I_d - \mathcal{S}^{m,g}_{w+a,d}))\epsilon_1(\mathbf{x},\mathbf{y};\boldsymbol{\theta})\|_{L^\infty(\Gamma^d\times\Gamma^d)}\\ &\leq \mathcal{Q}(\sigma,\delta^*(\sigma),d,\tilde{M}_{\epsilon_1})\eta^{\mu_3(\sigma,\delta^*(\sigma),d)}\exp\left(-\frac{d\sigma}{2^{1/d}}\eta^{\mu_2(d)}\right).\end{aligned}$$

where

$$\begin{aligned}\mathcal{Q}(\sigma,\delta^*(\sigma),d,\tilde{M}_{\epsilon_1}) = &\frac{8\mathcal{Q}(\sigma,\delta^*(\sigma),d,\tilde{M}(\phi))\eta^{\mu_3(\sigma,\delta^*(\sigma),d)}\exp\left(-\frac{d\sigma}{2^{1/d}}\eta^{\mu_2(d)}\right)C(\sigma)a(\delta,\sigma)}{\exp(\sigma\delta^*(\sigma)C_2(\sigma))e\delta\sigma}\\ &* \frac{\max\{1,C_1(\sigma,\delta^*(\sigma),d,\tilde{M}_{\epsilon_1})\}}{|1-C_1(\sigma,\delta^*(\sigma),d,\tilde{M}_{\epsilon_1})|}.\end{aligned}$$

It follows that

$$\left|\sum_{r=1}^N\sum_{h=1}^N \phi(\mathbf{x}_r,\mathbf{y}_h;\boldsymbol{\theta})\boldsymbol{\psi}^i_m[h]\boldsymbol{\psi}^j_q[r]\right| \leq \sqrt{n_m}\sqrt{n_q}\mathcal{M}(\sigma,\delta^*(\sigma),d,\tilde{M}(\phi))\eta^{2\mu_3(\sigma,\delta^*(\sigma),d)}\exp\left(-\frac{2d\sigma}{2^{1/d}}\eta^{\mu_2(d)}\right).$$

Now, if $w \leq d/\log(2)$ then from Theorem 2 for all $\mathbf{z}\in\mathcal{E}_{\hat{\sigma},\ldots,\hat{\sigma}}$

$$\|\epsilon_1(\cdot,\mathbf{z};\boldsymbol{\theta})\|_{L^\infty(\Gamma^d)} \leq \tilde{P}_{\epsilon_1}(\sigma,\delta^*(\sigma),d,\tilde{M}(\phi)).$$

From equation (11)

$$\begin{aligned}\|\epsilon_2(\mathbf{x},\mathbf{y};\boldsymbol{\theta})\|_{L^\infty(\Gamma^d\times\Gamma^d)} &\leq \frac{C_1(\sigma,\delta^*(\sigma),\tilde{P}_{\epsilon_1})}{|1-C_1(\sigma,\delta^*(\sigma),\tilde{P}_{\epsilon_1})|}\max\{1,C_1(\sigma,\delta^*(\sigma),\tilde{P}_{\epsilon_1})\}^d\eta^{-\mu_1}\\ &\leq \frac{8C(\sigma)a(\delta,\sigma)C_1(\sigma,\delta^*(\sigma),d,\tilde{M}(\phi))}{e\delta\sigma}\eta^{-2\mu_1}\\ &* \frac{\max\{1,C_1(\sigma,\delta^*(\sigma),d,\tilde{P}_{\epsilon_1})\}^d}{|1-C_1(\sigma,\delta^*(\sigma),d,\tilde{P}_{\epsilon_1})|}\frac{\max\{1,C_1(\sigma,\delta^*(\sigma),d,\tilde{M}(\phi)\}^d}{|1-C_1(\sigma,\delta^*(\sigma),d,\tilde{M}(\phi))|}.\end{aligned}$$

□

There are many cases where it is not possible, or hard to prove that the covariance function $\phi(\mathbf{x},\mathbf{y};\boldsymbol{\theta})$ has a bounded analytic extension on a well defined region in $\mathbb{C}^d\times\mathbb{C}^d$. For example, for the Matérn covariance function ($\nu = 0.5$), there is a derivative discontinuity at the origin. We look at a class of covariance functions where it is possible to prove decay rates of the covariance matrix $\mathbf{C}_W(\boldsymbol{\theta})$ whenever any two cells $B^i_m\in\mathcal{B}^i$ and $B^j_q\in\mathcal{B}^j$ are well separated.

Let $\mathbf{f}$ be a $\mathbb{R}^d$ valued vector, where for $n = 1,\ldots,d$, $f_n := \theta_n(x_n - y_n)^2$. The idea is to form a sparse grid representation of the covariance function $\phi(\mathbf{f}) : \mathbb{R}^d\to\mathbb{R}$ that avoids the analytic singularity at the origin.

Consider any two cells $B^i_m\in\mathcal{B}^i$ and $B^j_q\in\mathcal{B}^j$. Assume that the minimum Euclidean distance between any two observations $\mathbf{x}^*\in B^i_m$ and $\mathbf{y}^*\in B^j_q$ is $\tau_{i,j} > 0$ i.e. $\|\mathbf{x}^* - \mathbf{y}^*\| \geq \tau_{i,j}$. Conversely the maximal distance between all the observations in B^i_m and B^j_q is $\|\mathbf{x}^* - \mathbf{y}^*\| \leq 2\sqrt{d}$. Let $\{\mathbf{x}^{min},\mathbf{y}^{min}\}$ be the minimal distance pair observations and similarly, let $\{\mathbf{x}^{max},\mathbf{y}^{max}\}$ be the maximal. Thus for $n = 1,\ldots,d$, $f_n\in[\theta_n(x^{min}_n - y^{min}_n)^2,\theta_n(x^{max}_n - y^{max}_n)^2]$.

Let $\mathbf{t}$ be a $\mathbb{R}^d$ valued vector and for $n = 1, \ldots, d$ let

$$t_n := \frac{2(f_n - \theta_n(x_n^{min} - y_n^{min})^2)}{\theta_n((x_n^{max} - y_n^{max})^2 - (x_n^{min} - y_n^{min})^2)} - 1 = \frac{2((x_n - y_n)^2 - (x_n^{min} - y_n^{min})^2}{(x_n^{max} - y_n^{max})^2 - (x_n^{min} - y_n^{min})^2} - 1. \quad (19)$$

Thus $\mathbf{t} \in \Gamma^d$ and the covariance function is recast as $\phi(\mathbf{t}; \boldsymbol{\theta}) : \Gamma_d \to \mathbb{R}$.

Now, for this case the multi-level basis constructed from the observations $\mathbb{S}$ and the monomial set $\tilde{\mathcal{Q}}^a_{\Lambda^{m,g}(w)}$, for some $a, w \in \mathbb{N}_+$, is not enough to show decay of the coefficients of $\mathbf{C}_W(\boldsymbol{\theta})$. Instead, assume that the *extended* multi-level basis (see Remark 9) is constructed from the set of monomials $\tilde{\mathcal{Q}}^a_{\tilde{\Lambda}^{m,g}(w)}$, with extended Smolyak index set $\tilde{\Lambda}^{m,g}(w + a)$

Lemma 4. *Suppose that $\phi(\mathbf{t}; \boldsymbol{\theta}) \in C^0(\Gamma^d; \mathbb{R})$ has an analytic extension on $\mathcal{E}_{\hat{\sigma},\ldots,\hat{\sigma}}$ and is bounded by $\tilde{M}(\phi)$. Let $\mathcal{P}^{\tilde{p}}(\mathbb{S})^\perp$ be the subspace in $\mathbb{R}^N$ generated by the extended Smolyak index set $\tilde{\Lambda}^{m,g}(w+a)$ for some $w, a \in \mathbb{N}_+$. Let $\mathcal{S}^{m,g}_{w+a,d}$ be a sparse grid generated from the index set $\tilde{\Lambda}^{m,g}(w + a)$ with $\eta(w + a)$ collocation nodes. For $m, q = 0, \ldots, t$ consider any extended multi-level vector $\boldsymbol{\psi}^i_m \in \mathcal{P}^{\tilde{p}}(\mathbb{S})^\perp$, with n_m non-zero entries, from the cell B^i_m and any extended multi-level vector $\boldsymbol{\psi}^j_q \in \mathcal{P}^{\tilde{p}}(\mathbb{S})^\perp$, with n_q non-zero entries, from the cell B^j_q.*
(a) If $w + a > d/\log(2)$ then

$$\left| \sum_{r=1}^N \sum_{h=1}^N \phi(\mathbf{t}; \boldsymbol{\theta}) \boldsymbol{\psi}^i_m[h] \boldsymbol{\psi}^j_q[r] \right| \leq \sqrt{n_m n_q} \mathcal{Q}(\sigma, \delta^*(\sigma), d, \tilde{M}(\phi)) \eta^{\mu_3(\sigma,\delta^*(\sigma),d)} \exp\left(-\frac{d\sigma}{2^{1/d}} \eta^{\mu_2(d)} \right).$$

(b) If $w + a \leq d/\log(2)$ then

$$\left| \sum_{r=1}^N \sum_{h=1}^N \phi(\mathbf{t}; \boldsymbol{\theta}) \boldsymbol{\psi}^i_m[h] \boldsymbol{\psi}^j_q[r] \right| \leq \sqrt{n_m n_q} \frac{C_1(\sigma, \delta^*(\sigma), \tilde{M}(\phi))}{|1 - C_1(\sigma, \delta^*(\sigma), \tilde{M}(v))|} \max\left\{1, C_1(\sigma, \delta^*(\sigma), \tilde{M}(\phi))\right\}^d \eta^{-\mu_1}.$$

Proof. Now, perform the Smolyak sparse grid expansion in terms of the vector $\mathbf{t} \in \Gamma_d$ and let $\epsilon(\mathbf{t}; \boldsymbol{\theta}) := \phi(\mathbf{t}; \boldsymbol{\theta}) - \mathcal{S}^{m,g}_{w+a,d} \phi(\mathbf{t}; \boldsymbol{\theta})$.

$$\sum_{k=1}^N \sum_{l=1}^N \phi(\mathbf{t}(\mathbf{x}_k, \mathbf{y}_l); \boldsymbol{\theta}) \boldsymbol{\psi}^i_m[k] \boldsymbol{\psi}^j_q[l] = \sum_{k=1}^N \sum_{l=1}^N \left(\epsilon(\mathbf{t}(\mathbf{x}_k, \mathbf{y}_l); \boldsymbol{\theta}) + \sum_{\mathbf{i} \in \mathbb{N}^d_+ : g(\mathbf{i}) \leq w} \bigotimes_{n=1}^d \Delta_n^{m(i_n)}(\mathbf{t}(\mathbf{x}_l, \mathbf{y}_l); \boldsymbol{\theta}) \right) \boldsymbol{\psi}^i_m[k] \boldsymbol{\psi}^j_q[l].$$

Since $\boldsymbol{\psi}^i_m, \boldsymbol{\psi}^j_q \in \mathcal{P}^{\tilde{p}}(\mathbb{S})^\perp$

$$\sum_{k=1}^N \sum_{l=1}^N \left(\sum_{\mathbf{i} \in \mathbb{N}^d_+ : g(\mathbf{i}) \leq w} \bigotimes_{n=1}^d \Delta_n^{m(i_n)}(\mathbf{t}(\mathbf{x}_k, \mathbf{y}_l); \boldsymbol{\theta}) \right) \boldsymbol{\psi}^i_m[k] \boldsymbol{\psi}^j_q[l] = 0,$$

thus

$$\left|\sum_{k=1}^{N}\sum_{l=1}^{N}\phi(\mathbf{t}(\mathbf{x}_k,\mathbf{y}_l);\boldsymbol{\theta})\boldsymbol{\psi}_m^i[k]\boldsymbol{\psi}_q^j[l]\right| = \left|\sum_{k=1}^{N}\sum_{l=1}^{N}\epsilon(\mathbf{t}(\mathbf{x}_k,\mathbf{y}_l);\boldsymbol{\theta})\boldsymbol{\psi}_m^i[k]\boldsymbol{\psi}_q^j[l]\right|$$
$$\leq \|\epsilon(\mathbf{t};\boldsymbol{\theta})\|_{L^\infty(\Gamma^d)}\sum_{k=1}^{N}\sum_{l=1}^{N}|\boldsymbol{\psi}_m^i[k]||\boldsymbol{\psi}_q^j[l]|.$$

The result follows from Equations (10) and (11). □

7.2 Analyticity of the Matérn covariance function

In this section the analytic extension for the Matérn covariance function

$$\phi(r;\boldsymbol{\theta}) := \frac{1}{\Gamma(\nu)2^{\nu-1}}\left(\sqrt{2\nu}r(\boldsymbol{\theta})\right)^{\nu}K_\nu\left(\sqrt{2\nu}r(\boldsymbol{\theta})\right),$$

is analyzed on $\mathbb{C}^d$ or $\mathbb{C}^d\times\mathbb{C}^d$, where $r(\boldsymbol{\theta}) = \left(\mathbf{x}-\mathbf{y}\right)^T diag(\boldsymbol{\theta})(\mathbf{x}-\mathbf{y})\big)^{\frac{1}{2}}$, $\boldsymbol{\theta} = [\theta_1,\ldots,\theta_d]\in\mathbb{R}_+^d$ are positive constants, $diag(\theta)\in\mathbb{R}^{d\times d}$ is a zero matrix with the vector $\boldsymbol{\theta}$ on the diagonal, and $\mathbf{x},\mathbf{y}\in\Gamma_d$.

Recall from Remark 6 that for different values of $\nu\in\mathbb{R}_+$ different shapes of the covariance function are obtained, i.e. for $\nu = 1/2+n$, $n\in\mathbb{N}_+$

$$\phi(r;\nu,\boldsymbol{\theta}) = \exp(-\sqrt{2\nu}r(\boldsymbol{\theta}))\frac{\Gamma(n+1)}{\Gamma(2n+1)}\sum_{k=1}^{n}\frac{(n+1)!}{n!(n-k)!}(\sqrt{8v}r(\boldsymbol{\theta}))^{n-k}.$$

The Gaussian covariance function is first analyzed, which is the case for the Matérn covariance function when $\nu\to\infty$. However, the covariance function is recast in a slightly different form:

$$\phi(\mathbf{x},\mathbf{y};\theta) = \exp\left(-(\mathbf{x}-\mathbf{y})^T diag(\boldsymbol{\theta})(\mathbf{x}-\mathbf{y})\right).$$

Lemma 5. *There exists an analytic extension of the Gaussian correlation function* $\phi(\mathbf{x},\mathbf{y};\boldsymbol{\theta}) : \Gamma_d\times\Gamma_d\to\mathbb{R}$ *on* $(\mathbf{x},\mathbf{y})\in\prod_{n=1}^d\mathcal{E}_{n,\hat{\sigma}}\times\prod_{n=1}^d\mathcal{E}_{n,\hat{\sigma}}\subset\mathbb{C}^d$. *Furthermore, for all* $(\mathbf{x},\mathbf{y})\in\prod_{n=1}^d\mathcal{E}_{n,\hat{\sigma}}\times\prod_{n=1}^d\mathcal{E}_{n,\hat{\sigma}}$

$$|\phi(\mathbf{x},\mathbf{y};\boldsymbol{\theta})|\leq\exp\left(2d(e^{2\hat{\sigma}}+e^{-2\hat{\sigma}})\right)\prod_{n=1}^{d}\exp(\theta_n).$$

Proof. First, the following facts from complex analysis are stated:

i) The exponential function is holomorphic on $\mathbb{C}$.

ii) The polynomial functions are holomorphic on $\mathbb{C}$.

iii) The sum of two holomorphic functions on $\mathbb{C}$ are holomorphic on $\mathbb{C}$.

iv) The composition of two holomorphic functions on $\mathbb{C}$ are holomorphic on $\mathbb{C}$.

From *ii*) and *iii*) for all $\mathbf{x}, \mathbf{y} \in \mathbb{C}^d$ it follows that $(\mathbf{x}-\mathbf{y})^T diag(\boldsymbol{\theta})(\mathbf{x}-\mathbf{y}) : \mathbb{C}^d \times \mathbb{C}^d \to \mathbb{C}$ is holomorphic. From *i*) and *iv*) it is concluded that $\phi(\mathbf{x}, \mathbf{y}; \boldsymbol{\theta}) : \mathbb{C}^d \times \mathbb{C}^d \to \mathbb{C}$ is analytic.

The next step is to bound the function $\phi(\mathbf{x}, \mathbf{y}; \boldsymbol{\theta}) : \mathbb{C}^d \times \mathbb{C}^d \to \mathbb{C}$. However, for the isotropic sparse grid representation used in this paper it is sufficient to show that the analytic extension of $\phi(\mathbf{x}, \mathbf{y}; \boldsymbol{\theta})$ is valid on $\prod_{n=1}^d \mathcal{E}_{n,\hat{\sigma}} \times \prod_{n=1}^d \mathcal{E}_{n,\hat{\sigma}}$.

Now, recall the shape of the polyellipse

$$\mathcal{E}_{n,\hat{\sigma}_n} = \left\{ z \in \mathbb{C}, \hat{\sigma} \geq \delta_n \geq 0;\, \mathrm{Re}\, z = \frac{e^{\delta_n} + e^{-\delta_n}}{2} cos(\theta),\;\; \mathrm{Im}\, z = \frac{e^{\delta_n} - e^{-\delta_n}}{2} sin(\theta), \theta \in [0, 2\pi) \right\}.$$

Therefore any $z \in \mathcal{E}_{n,\hat{\sigma}}$ can be written in the form $\frac{(e^{2\delta_n}+e^{-2\delta_n})^{1/2}}{2} e^{i\Omega_n}$, $i = \sqrt{-1}$, $\hat{\sigma} \geq \delta_n \geq 0$ and $\Omega_n \in \mathbb{R}$. Thus for all $\mathbf{x}, \mathbf{y} \in \prod_{n=1}^d \mathcal{E}_{n,\hat{\sigma}}$

$$\begin{aligned} |\phi(\mathbf{x}, \mathbf{y}; \boldsymbol{\theta})| &\leq \exp\left(\sum_{n=1}^d \left| \theta_n (x_n - y_n)^2 \right| \right) \leq \exp\left(\sum_{n=1}^d \theta_n 2(e^{2\hat{\sigma}} + e^{-2\hat{\sigma}}) \right) \\ &\leq \prod_{n=1}^d \exp(\theta_n)^{2d(e^{2\hat{\sigma}} + e^{-2\hat{\sigma}}))}. \end{aligned}$$

□

Now, the analyticity of the Matérn covariance function

$$\phi(r; \boldsymbol{\theta}) := \frac{1}{\Gamma(\nu) 2^{\nu-1}} \left(\sqrt{2\nu} r(\boldsymbol{\theta}) \right)^{\nu} K_\nu \left(\sqrt{2\nu} r(\boldsymbol{\theta}) \right).$$

is analyzed. Since, in general, the Matérn covariance function is not differentiable at the origin then there is no analytic extension on $\prod_{n=1}^d \mathcal{E}_{n,\hat{\sigma}} \times \prod_{n=1}^d \mathcal{E}_{n,\hat{\sigma}}$, thus Lemma 4 will not be valid. However, it is not necessary to show that an analytic extension of $\phi(\mathbf{x}, \mathbf{y}; \boldsymbol{\theta}) : \Gamma^d \times \Gamma^d \to \mathbb{R}$ exists for the entire domain $\Gamma^d \times \Gamma^d$, but only on a sub-domain that does not cross the origin.

Consider any two cells $B_m^i \in \mathcal{B}^i$ and $B_q^j \in \mathcal{B}^j$, at levels i and j, and assume that the minimum Euclidean distance between any two observations $\mathbf{x}^* \in B_m^i$ and $\mathbf{y}^* \in B_q^j$ is $\tau_{i,j} > 0$ i.e. $\|\mathbf{x}^* - \mathbf{y}^*\| \geq \tau_{i,j}$. Conversely the maximal distance between all observations in B_m^i and B_q^j is $\|\mathbf{x}^* - \mathbf{y}^*\| \leq 2\sqrt{d}$. Let $\{\mathbf{x}^{min}, \mathbf{y}^{min}\}$ be the minimal distance pair observations and similarly, let $\{\mathbf{x}^{max}, \mathbf{y}^{max}\}$ be the maximal. Furthermore, for $n = 1, \ldots, d$, let $m_n := x_n^{min} - y_n^{min}$ and $M_n := x_n^{max} - y_n^{max}$.

Let $\mathbf{t} \in \mathbb{R}^d$, where for $n = 1, \ldots, d$, t_n is equal to equation (19), and recast the Matérn covariance function as $\phi(\mathbf{t}; \nu, \theta) : \Gamma_d \to \mathbb{R}$ where

$$\phi(\mathbf{t}; \nu, \boldsymbol{\theta}) := \frac{1}{\Gamma(\nu) 2^{\nu-1}} \left(\sqrt{2\nu} r(\mathbf{t}, \boldsymbol{\theta}) \right)^{\nu} K_\nu \left(\sqrt{2\nu} r(\mathbf{t}, \boldsymbol{\theta}) \right),$$

$\mathbf{t} \in \Gamma^d$ and $r(\mathbf{t}, \boldsymbol{\theta}) = \sum_{n=1}^d \theta_n ((t_n + 1)(M_n^2 - m_n^2)/2 + m_n^2)$.

The next step is to find a suitable value for $\hat{\sigma} > 0$ such that the Matérn function is analytic and bounded on the region $\prod_{n=1}^d \mathcal{E}_{n,\hat{\sigma}}$. In the following Lemma an estimate of the region of analyticity is obtained.

Lemma 6. *Given any two cells $B_m^i \in \mathcal{B}^i$ and $B_q^j \in \mathcal{B}^j$, such that the minimum Euclidean distance between any two observations $\mathbf{x}^* \in B_m^i$ and $\mathbf{y}^* \in B_q^j$ is $\tau_{i,j} > 0$, then there exists an analytic*

extension of $\phi(\mathbf{t};\nu,\boldsymbol{\theta}):\Gamma_d\to\mathbb{R}$ *on* $\prod_{n=1}^d \mathcal{E}_{n,\hat{\sigma}_{i,j}}\subset\mathbb{C}^d$ *where*

$$\hat{\sigma}_{i,j} < \log\left(\frac{\tau_{i,j}}{2d}+1-\sqrt{\frac{\tau_{i,j}}{2d}\left(\frac{\tau_{i,j}}{2d}+2\right)}\right).$$

Moreover, for all $\mathbf{t}\in\prod_{n=1}^d \mathcal{E}_{n,\hat{\sigma}}\subset\mathbb{C}^d$ *it follows that* $|r(\mathbf{t};\boldsymbol{\theta})|\le\alpha(\boldsymbol{\theta},\hat{\sigma}_{i,j}):=\|\boldsymbol{\theta}\|_2\left(2(3+\sqrt{2}e^{\hat{\sigma}_{i,j}})\right)^{\frac{1}{2}}$ *and for* $n\in\mathbb{N}_+$

$$|\phi(r;1/2+n,\boldsymbol{\theta})|\le\exp(-(4(3+\sqrt{2}))^{1/2}\nu\|\boldsymbol{\theta}\|_2)\frac{\Gamma(n+1)}{\Gamma(2n+1)}\sum_{k=1}^{n}\frac{(n+1)!}{n!(n-k)!}(\sqrt{8v}\alpha(\boldsymbol{\theta},\hat{\sigma}_{i,j}))^{n-k}.$$

Proof. From equation (18) for $n=1,\dots,d$

$$f_n=\theta_n\left(m_n^2+\frac{(t_n+1)(M_n^2-m_n^2)}{2}\right),$$

thus, $\sum_{n=1}^d m_n^2\ge\tau_{i,j}$ and for $n=1,\dots,d$, $m_n^2\le M_n^2\le 4$. Recall that the Matérn function can be rewritten as

$$\phi(\mathbf{f};\nu,\boldsymbol{\theta})=\frac{1}{\Gamma(\nu)2^{\nu-1}}\left(\sqrt{2\nu}r(\mathbf{f},\boldsymbol{\theta})\right)^{\nu}K_\nu\left(\sqrt{2\nu}r(\mathbf{f},\boldsymbol{\theta})\right),$$

where $r(\mathbf{f},\boldsymbol{\theta})=\left(\sum_{n=1}^d f_n\right)^{\frac{1}{2}}$. Now, the square root function $z^{\frac{1}{2}}$, $z\in\mathbb{C}$, is analytic if $\operatorname{Re}z>0$. Given that the modified Bessel functions are analytic in the complex plane and polynomials are analytic in $\mathbb{C}$, then $\phi(\mathbf{t};\nu\boldsymbol{\theta})$ is analytic if $\mathbf{t}$ belongs in a region in $\mathbb{C}^d$ such that

$$\operatorname{Re}\sum_{n=1}^d f_n\ge\theta_{min}\sum_{n=1}^d\operatorname{Re}\left(m_n^2+\frac{(t_n+1)(M_n^2-m_n^2)}{2}\right)>0.$$

This condition is satisfied if for $n=1,\dots,d$, $\operatorname{Re}t_n>-1-\frac{\tau_{i,j}}{2d}$, i.e. $\operatorname{Re}\mathbf{t}\in\prod_{n=1}^d(-1-\frac{\tau_{i,j}}{2d},\infty)$.

The next step is to embed a polyellipse $\mathcal{E}_{\log\rho}\subset(-1-\frac{\tau_{i,j}}{2d},\infty)\times(-\infty i,\infty i)$, as shown in Figure 6. This is achieved by setting $\rho<\left(\frac{\tau_{i,j}}{2d}+1-\sqrt{\frac{\tau_{i,j}}{2d}\left(\frac{\tau_{i,j}}{2d}+2\right)}\right)$.

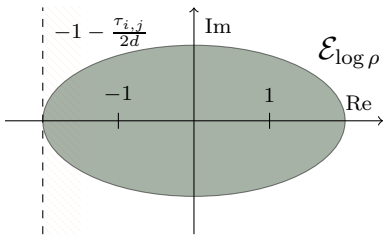

Figure 6: Embedding of polyellipse $\mathcal{E}_{\log\rho}$ in the analyticity region $\operatorname{Re}z>-1-\frac{\tau_{i,j}}{2d}$. The embedding is achieved for a polyellipse with $\rho<\left(\frac{\tau_{i,j}}{2d}+1-\sqrt{\frac{\tau_{i,j}}{2d}\left(\frac{\tau_{i,j}}{2d}+2\right)}\right)$.

Now,

$$\max_{t\in\mathcal{E}_{\log\rho}}|t|=\left(\left(\frac{\rho+\rho^{-1}}{2}\right)^2+\left(\frac{\rho-\rho^{-1}}{2}\right)^2\right)^{1/2}\le\sqrt{2}\rho,$$

then for $\mathbf{t} \in \prod_{n=1}^{d} \mathcal{E}_{n,\hat{\sigma}}$

$$|r(\mathbf{t};\boldsymbol{\theta})| \leq \left(\sum_{n=1}^{d} \theta_n^2 \left(m_n^2 + \frac{(|t_n|+1)(M_n^2 - m_n^2)}{2}\right)\right)^{1/2} \leq \left(\sum_{n=1}^{d} \theta_n^2 \left(4 + (|t_n|+1)4\right)\right)^{1/2}$$
$$\leq \|\boldsymbol{\theta}\|_2 \left(2(3+\sqrt{2}e^{\hat{\sigma}})\right)^{\frac{1}{2}}.$$

The result follows. □

7.3 A-posteriori error estimates

Given that the region of analyticity of the Matérn covariance function and decay rates of the covariance matrix haven been derived, estimates for the error between the full $\mathbf{x}_W = \mathbf{C}_W^n(\boldsymbol{\theta})^{-1}\mathbf{Z}_W^n$ and sparse solution $\tilde{\mathbf{x}}_W = \tilde{\mathbf{C}}_W^n(\boldsymbol{\theta})^{-1}\mathbf{Z}_W^n$ are shown in this section. For $n = t, \ldots, -1$ suppose that $\sigma_{min}(\mathbf{C}_W^n(\boldsymbol{\theta}))\|\mathbf{E}\|_2 < 1$, where $\mathbf{E} := \mathbf{C}_W^n(\boldsymbol{\theta}) - \tilde{\mathbf{C}}_W^n(\boldsymbol{\theta})$, then

$$\begin{aligned}
\|\mathbf{C}_W^n(\boldsymbol{\theta})^{-1} - \tilde{\mathbf{C}}_W^n(\boldsymbol{\theta})^{-1}\|_2 &= \|(\tilde{\mathbf{C}}_W^n(\boldsymbol{\theta}) + \mathbf{E})^{-1} - \tilde{\mathbf{C}}_W^n(\boldsymbol{\theta})^{-1}\|_2 \\
&= \|\tilde{\mathbf{C}}_W^n(\boldsymbol{\theta})^{-1}\mathbf{E}(\mathbf{I} - \tilde{\mathbf{C}}_W^n(\boldsymbol{\theta})^{-1}\mathbf{E})^{-1}\tilde{\mathbf{C}}_W^n(\boldsymbol{\theta})^{-1}\|_2 \\
&\leq \sigma_{min}^{-2}(\tilde{\mathbf{C}}_W^n(\boldsymbol{\theta}))\|(\mathbf{I} - \tilde{\mathbf{C}}_W^n(\boldsymbol{\theta})^{-1}\mathbf{E})^{-1}\|_2\|\mathbf{E}\|_2 \\
&\leq \sigma_{min}^{-2}(\tilde{\mathbf{C}}_W^n(\boldsymbol{\theta}))\sigma_{min}(\mathbf{I} - \tilde{\mathbf{C}}_W^n(\boldsymbol{\theta})^{-1}\mathbf{E})\|\mathbf{E}\|_2 \\
&\leq \sigma_{min}^{-2}(\tilde{\mathbf{C}}_W^n(\boldsymbol{\theta}))(1 + \sigma_{min}(\tilde{\mathbf{C}}_W^n(\boldsymbol{\theta})^{-1}\mathbf{E}))\|\mathbf{E}\|_2 \\
&\leq \sigma_{min}^{-2}(\tilde{\mathbf{C}}_W^n(\boldsymbol{\theta}))(1 + \sigma_{min}(\tilde{\mathbf{C}}_W^n(\boldsymbol{\theta})^{-1}\mathbf{E}))\|\mathbf{E}\|_2 \\
&\leq \sigma_{min}^{-2}(\tilde{\mathbf{C}}_W^n(\boldsymbol{\theta}))(1 + \sigma_{min}(\tilde{\mathbf{C}}_W^n(\boldsymbol{\theta})^{-1})\|\mathbf{E}\|_2)\|\mathbf{E}\|_2 \\
&\leq \sigma_{min}^{-2}(\tilde{\mathbf{C}}_W^n(\boldsymbol{\theta}))(1 + \sigma_{max}^{-1}(\tilde{\mathbf{C}}_W^n(\boldsymbol{\theta}))\|\mathbf{E}\|_2)\|\mathbf{E}\|_2.
\end{aligned}$$

Note that the following singular value inequalities, which can be found in [18], where applied: If $A, B \in \mathbb{C}^{n,n}$ then

i) $\sigma_{min}(1 + A) \leq 1 + \sigma_{min}(A)$.

ii) $\sigma_{max}(AB) \leq \sigma_{max}(A)\sigma_{max}(B)$.

iii) $\sigma_{min}(AB) \leq \sigma_{min}(A)\sigma_{max}(B)$.

iv) $\sigma_{max}(A^{-1}) = \sigma_{min}^{-1}(A)$.

Remark 13. *The terms $\sigma_{min}(\tilde{\mathbf{C}}_W^n(\boldsymbol{\theta}))$ and $\sigma_{max}(\tilde{\mathbf{C}}_W^n(\boldsymbol{\theta}))$ can be estimated by using a Lanzcos sparse Singular Value Decomposition (SVD) type algorithm [5, 20].*

The term $\|\mathbf{E}\|_2$ can be bounded as

$$\|\mathbf{E}\|_2 \leq \sum_{i=\max\{n,1\}}^{t} \sum_{j=\max\{n,1\}}^{t} \|\mathbf{E}^{i,j}\|_2,$$

where for $i, j = \max\{n, 1\}, \ldots, t$, $\mathbf{E}^{i,j} := \mathbf{C}_W^{i,j}(\boldsymbol{\theta}) - \tilde{\mathbf{C}}_W^{i,j}(\boldsymbol{\theta})$.

Remark 14. *It is not too hard to see that for $\tau_{i,j} \geq 0$ Algorithm 5 leads to $\mathbf{E}^{i,j} = \mathbf{0}$ if $i = -1, 0$ or $j = -1, 0$.*

The contributions of the errors $\|\mathbf{E}^{i,j}\|_2$, for $i, j = \max\{n, 1\}, \ldots, t$, can now be estimated, but first, the following assumption is made:

Assumption 2. *For $i, j, = 1, \ldots, t$ let $\tau_{i,j} := \tau 2^{t-(i+j)/2}$, for some user defined parameter $\tau > 0$.*

Now, for $i, j = \max\{1, n\}, \ldots, t$ if $w + a > d/\log(2)$ then from Assumption 2 and Lemma 1, 2 and 3

$$\begin{aligned}\|\mathbf{E}^{i,j}\|_2 &\leq \sqrt{2^{i+j}}\tilde{p}\max\{|\mathbf{E}^{i,j}(:)|\} \\ &\leq 2^{t+1}\tilde{p}^2\mathcal{M}(\sigma, \delta^*(\sigma), d, \tilde{M}(\phi))\eta^{2\mu_3(\sigma,\delta^*(\sigma),d)}\exp\left(-\frac{2d\sigma}{2^{1/d}}\eta(w+a)^{\mu_2(d)}\right),\end{aligned}$$

and thus

$$\begin{aligned}\|\mathbf{E}\|_2 \leq 2^{t+1}\tilde{p}^2(t - \max\{n, 1\})^2\mathcal{M}(\sigma, \delta^*(\sigma), d, \tilde{M}(\phi)) \\ \eta(w+a)^{2\mu_3(\sigma,\delta^*(\sigma),d)}\exp\left(-\frac{2d\sigma}{2^{1/d}}\eta(w+a)^{\mu_2(d)}\right).\end{aligned}$$

Following a similar argument for the case that $w + a \leq d/\log(2)$ the following theorem is proved:

Theorem 7. *Suppose the same conditions as in Lemma 3 are satisfied, then from Assumption 2*

$$\|\mathbf{C}_W^n(\boldsymbol{\theta})^{-1} - \tilde{\mathbf{C}}_W^n(\boldsymbol{\theta})^{-1}\|_2 \leq \sigma_{min}^{-2}(\tilde{\mathbf{C}}_W^n(\boldsymbol{\theta}))(1 + \sigma_{max}^{-1}(\tilde{\mathbf{C}}_W^n(\boldsymbol{\theta}))\|\mathbf{E}\|_2)\|\mathbf{E}\|_2.$$

where: (a) If $w + a > d/\log(2)$ then

$$\begin{aligned}\|\mathbf{E}\|_2 \leq 2^{t+1}\tilde{p}^2(t - \max\{n, 1\})^2\mathcal{M}(\sigma, \delta^*(\sigma), d, \tilde{M}(\phi)) \\ \eta(w+a)^{2\mu_3(\sigma,\delta^*(\sigma),d)}\exp\left(-\frac{2d\sigma}{2^{1/d}}\eta(w+a)^{\mu_2(d)}\right).\end{aligned}$$

(b) Else if $w + a \leq d/\log(2)$ then

$$\begin{aligned}\|\mathbf{E}\|_2 \leq 2^{t+4}\tilde{p}^2(t - \max\{n, 1\})^2\frac{C(\sigma)a(\delta, \sigma)C_1(\sigma, \delta^*(\sigma), d, \tilde{M}(\phi))}{e\delta\sigma}\eta^{-2\mu_1} \\ \frac{\max\{1, C_1(\sigma, \delta^*(\sigma), d, \tilde{P}_{\epsilon_1})\}^d}{|1 - C_1(\sigma, \delta^*(\sigma), d, \tilde{P}_{\epsilon_1})|}\frac{\max\{1, C_1(\sigma, \delta^*(\sigma), d, \tilde{M}(\phi)\}^d}{|1 - C_1(\sigma, \delta^*(\sigma), d, \tilde{M}(\phi))|}.\end{aligned}$$

Remark 15. *By combining Theorem 7 and Lemma 5 the a-posteriori error bound is computed for the Gaussian covariance function. For this bound the user can set the size of the polyellipse parameter $\hat{\sigma}$ to one, thus $\sigma = 1/2$, $\tilde{M}(\phi) = \prod_{n=1}^{d}\exp(\theta_n)^{2d(e^{2\hat{\sigma}}+e^{-2\hat{\sigma}}))}$.*

Now one must deal with the more general Matérn covariance function error estimates . It is straightforward to prove the following theorem.

Theorem 8. *Suppose the same conditions as in Lemma 4 and Lemma 6 are satisfied and assume that $\sigma_{i,j} := \hat{\sigma}_{i,j}/2$,*

$$0 < \hat{\sigma}_{i,j} < \log\left(\frac{\tau_{i,j}}{2d} + 1 - \sqrt{\frac{\tau_{i,j}}{2d}\left(\frac{\tau_{i,j}}{2d} + 2\right)}\right),$$

for $i, j = 0, \ldots, t$, *and*

$$\tilde{M}(\phi) := \exp(-(4(3+\sqrt{2}))^{1/2}\nu\|\boldsymbol{\theta}\|_2)\frac{\Gamma(n+1)}{\Gamma(2n+1)}\sum_{k=1}^{n}\frac{(n+1)!}{n!(n-k)!}(\sqrt{8v}\alpha(\boldsymbol{\theta},\hat{\sigma}_{i,j}))^{n-k}$$

for $n = -1, \ldots, t$. *Furthermore, assume the sparse covariance matrix* $\tilde{\mathbf{C}}_W^n(\boldsymbol{\theta})$ *non zero entries are computed with Algorithm 7 then for* $n = -1, \ldots, t$,

$$\|\mathbf{C}_W^n(\boldsymbol{\theta})^{-1} - \tilde{\mathbf{C}}_W^n(\boldsymbol{\theta})^{-1}\|_2 \leq \sigma_{min}^{-2}(\tilde{\mathbf{C}}_W^n(\boldsymbol{\theta}))(1+\sigma_{max}^{-1}(\tilde{\mathbf{C}}_W^n(\boldsymbol{\theta}))\|\mathbf{E}\|_2)\|\mathbf{E}\|_2.$$

where: (a) If $w + a > d/\log(2)$ *then*

$$\|\mathbf{E}\|_2 \leq 2^{t+1}\tilde{p}^2\sum_{i=\max\{n,1\}}^{t}\sum_{j=\max\{n,1\}}^{t}\mathcal{Q}(\sigma_{i,j},\delta^*(\sigma_{i,j}),d,\tilde{M}(\phi))\eta(w+a)^{\mu_3(\sigma,\delta^*(\sigma_{i,j}),d)}$$
$$\exp\left(-\frac{d\sigma_{i,j}}{2^{1/d}}\eta(w+a)^{\mu_2(d)}\right)$$

else if $w + a \leq d/\log(2)$ *then*

$$\|\mathbf{E}\|_2 \leq 2^{t+1}\tilde{p}^2\sum_{i=\max\{n,1\}}^{t}\sum_{j=\max\{n,1\}}^{t}\frac{C_1(\sigma_{i,j},\delta^*(\sigma_{i,j}),\tilde{M}(\phi))}{|1-C_1(\sigma_{i,j},\delta^*(\sigma_{i,j}),\tilde{M}(\phi))|}\max\left\{1,C_1(\sigma_{i,j},\delta^*(\sigma_{i,j}),\tilde{M}(\phi))\right\}^d$$
$$*\,\eta(w+a)^{-\mu_1(\sigma_{i,j})}.$$

8 Numerical results

The performance of the multi-level solver for estimation and prediction formed from random datasets in high dimensions is tested. The results show that the computational burden is significantly reduced. In particular, it is possible to now solve ill-conditioned problems efficiently. The implementation of the code is done with a combination of MATLAB and C++ packages:

i) **Matlab:** The RP/kD binary tree, multi-level basis construction, formation of the sparse matrix $\mathbf{C}_W$, estimation and prediction components are written and executed on Matlab [23]. The computational bottlenecks are executed on the C++ software packages, dynamic shared libraries, and the highly optimized BLAS and LAPACK packages contained in MATLAB.

ii) **Direct and fast summation:** The matlab code estimates the computational cost between the direct and fast summation methods and chooses the most efficient approach. For the direct method the highly optimized MATLAB arithmetic functions are used. For the fast summation method the KIFMM ($d = 3$) or ASKIT ($d > 3$) c++ codes are used. The KIFMM is modified to include a Hermite interpolant approximation of the Matérn covariance function, which is implemented with the intel MKL package [1] (see [8] for details). The ASKIT code is only used for the Gaussian case since the performance for the Matérn kernel is only slightly better than the direct approach.

iii) **Dynamic shared libraries:** These are produced with the GNU gcc/g++ packages. These libraries implement the Hermite interpolant with the intel MKL package (about 10 times faster than Matlab Matérn interpolant) and link the MATLAB code to the KIFMM.

iv) **Cholesky and determinant computation:** The Suite Sparse 4.2.1 package ([11, 13–16]) is used for the determinant computation of the sparse matrix $\tilde{\mathbf{C}}_W(\boldsymbol{\theta})$.

The code is executed on a single CPU (4 core Intel i7-3770 CPU @ 3.40GHz.) with Linux Ubuntu 14.04 and 32 GB memory. To test the effectiveness of the Multi-Level solver the following data sets are generated:

a) **Random hypercube data set:** The set of random observation locations $\mathbf{C}_0^d, \ldots, \mathbf{C}_{10}^d$ vary from 1,000, 2000, 4,000 to 512,000 knots generated on the hypercube $[-1, 1]^d$ for d dimensions. The observations locations are also nested, i.e. $\mathbf{C}_0^d \subset \cdots \subset \mathbf{C}_{10}^d$.

b) **Random n-sphere data set:** The set of nested random observations $\mathbf{S}_0^d \subset \cdots \subset \mathbf{S}_{10}^d$ varies from 1,000, 2000, 4000 to 128,000 knots generated on the n-sphere $\mathbf{S}_{d-1} := \{\mathbf{x} \in \mathbb{R}^d \mid \|\mathbf{x}\|_2 = 1\}$.

c) **Gaussian data set** The set of observations values $\mathbf{Z}_0^d, \mathbf{Z}_1^d, \ldots \mathbf{Z}_5^d$ are formed from the Gaussian random field model (1) for 1,000, 2,000, ... $32,000$ observation locations. The data set $\mathbf{Z}_n^d$ is generated from the set of nodes $\mathbf{C}_n^d$ or $\mathbf{S}_n^d$, with the covariance parameters (ν, ρ) and the corresponding set of monomials $\mathcal{Q}_{\Lambda^{m,g}(w)}$. Due to memory limitations 32,000 was the maximum size of the generated data set.

d) All the numerical test are done assuming the $\tilde{p} = p$.

8.1 Numerical stability and sparsity of the covariance multi-level matrix $\mathbf{C}_W$

The covariance matrix $\mathbf{C}(\boldsymbol{\theta})$ becomes increasingly ill-conditioned as the number of observations are increased, leading to instability of the numerical solver. It is now shown how effective Theorem 3 becomes in practice. In Figure 7 the condition number of the multi-level covariance matrix $\mathbf{C}_W(\boldsymbol{\theta})$ is plotted with respect to the cardinality p of the index set $\Lambda(w)$ for different w levels. The multilevel covariance matrix $\mathbf{C}_W(\boldsymbol{\theta})$ is built from the random cube $\mathbf{C}_4^d$ or n-sphere $\mathbf{S}_4^d$ observations. The covariance function is set to Matérn with $\nu = 1$ and $\rho = 1, 10, 100$. As the plots confirm the covariance matrix condition number significantly improves with increasing level w for the TD, SM and HC index sets. This is in contrast with the large condition numbers of the original covariance matrix $\mathbf{C}(\boldsymbol{\theta})$.

In Figure 8(a) the magnitude of the multi-level covariance matrix $\mathbf{C}_W(\boldsymbol{\theta})$ is plotted for $N = 8,000$ observations from the the n-sphere $\mathbf{S}_3^3$ with Matérn covariance parameters $\nu = 0.5$ and $\rho = 10$. Due to the large value of ρ the overlap between the covariance function at the different locations in $\mathbf{S}_3^3$ is quite significant, thus leading to a dense covariance matrix $\mathbf{C}(\boldsymbol{\theta})$ where the coefficients decay slowly. This is in contrast to the large number of small entries for $\mathbf{C}_W(\boldsymbol{\theta})$, as shown in the histogram in Figure 8(b). Note that the histogram is in terms of $\log_{10}$ of the absolute value of the entries of $\mathbf{C}_W(\boldsymbol{\theta})$. From the histogram it is observed that almost all the entries are more than 1000 smaller than the largest entries.

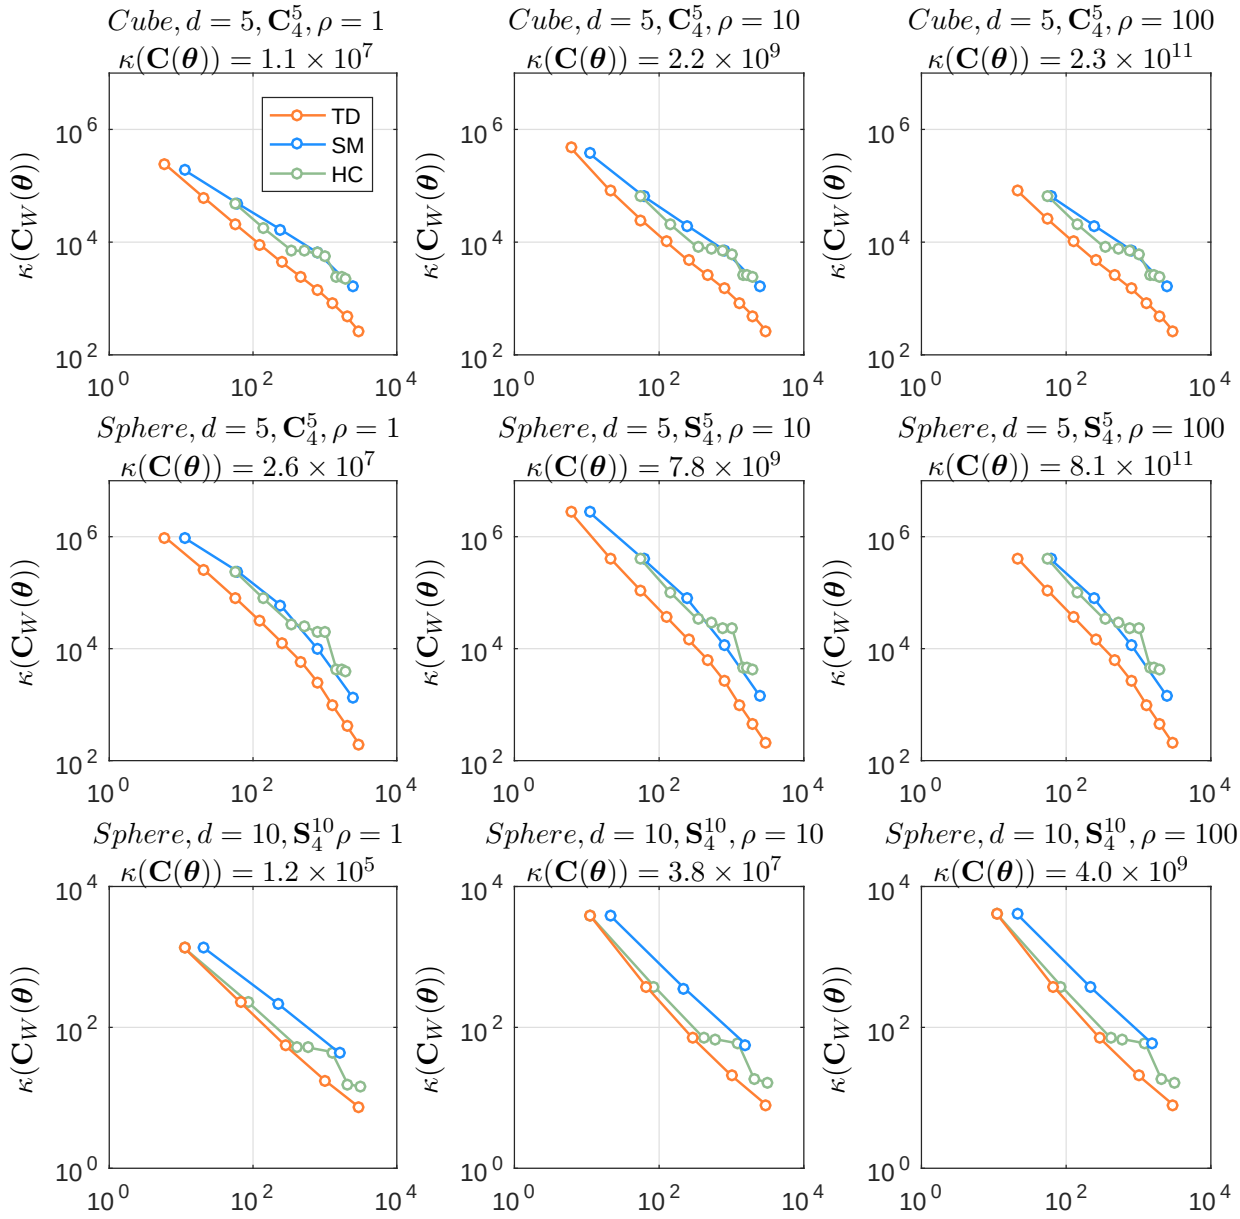

Figure 7: Condition number of the multi-level covariance matrix $\mathbf{C}_W(\boldsymbol{\theta})$ with respect to the size p of the polynomial approximation space with Total Degree, Smolyak and Hyperbolic index sets $\Lambda(w)$. The number of observations corresponds to 16,000 nodes generated on a hypercube or hyper-sphere of dimension $d = 5$ or $d = 10$. The covariance function is chosen to be Matérn with $\nu = 1$ and $\rho = 1, 10, 100$. The condition number of the covariance matrix $\mathbf{C}(\boldsymbol{\theta})$ is placed on the top of each subplot. The MB is constructed from a RP tree. As expected, as p increases with w the condition number of $\mathbf{C}_W(\boldsymbol{\theta})$ decreases significantly.

In Figure 8(a) and (b) the log determinant relative error $\frac{|\log\det\tilde{\mathbf{C}}_W(\boldsymbol{\theta}) - \log\det\mathbf{C}_W(\boldsymbol{\theta})|}{|\log\det\mathbf{C}_W(\boldsymbol{\theta})|}$ is plotted with respect to the sparsity of $\tilde{\mathbf{C}}_W(\boldsymbol{\theta})$ by choosing an increasingly larger distance criterion parameter τ. As it is observed even for highly sparse multi-level covariance matrices the relative error is about 1% and further improves as the sparsity is increased. Note that sparsity equal to one means that the matrix is fully populated.

In Table 3 sparsity and construction wall clock times of the sparse matrices $\tilde{\mathbf{C}}_W^i(\boldsymbol{\theta})$, $i =$

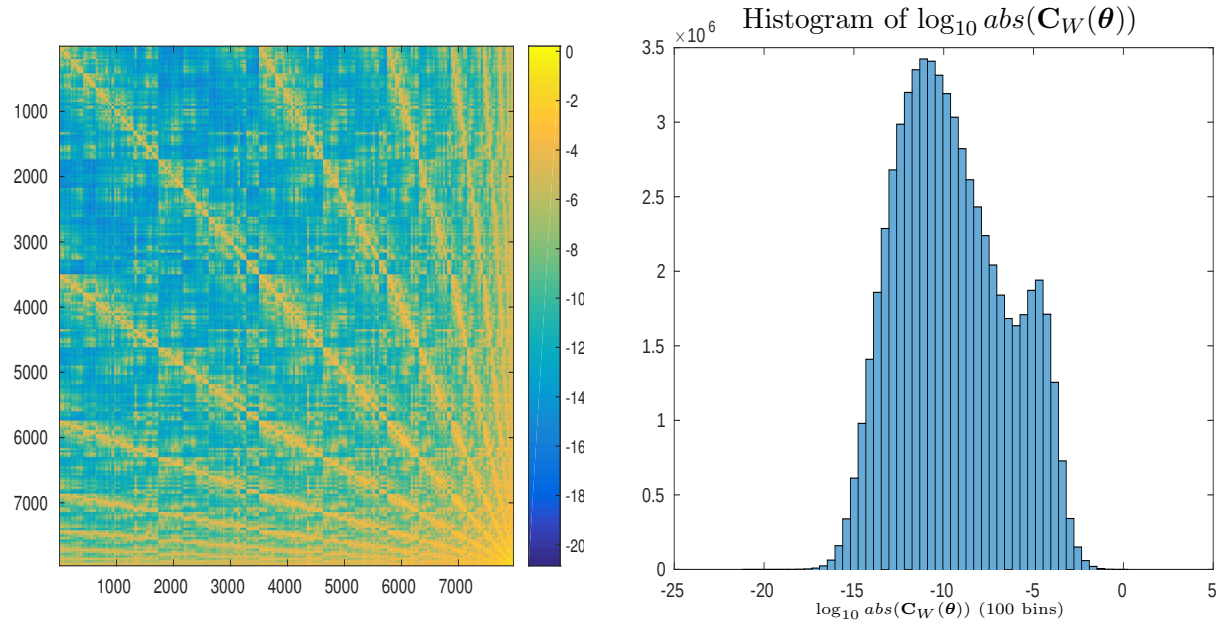

(a) (b)

Figure 8: (a) Magnitude pattern and (b) histogram of $\log_{10} abs(\mathbf{C}_W(\boldsymbol{\theta}))$ with 100 bins where $abs(\mathbf{C}_W(\boldsymbol{\theta})) \in \mathbb{R}^{(N-p)\times(N-p)}$ is the magnitude of the entries of the matrix $\mathbf{C}_W(\boldsymbol{\theta})$. The matrix $\mathbf{C}_W(\boldsymbol{\theta})$ is created with $d = 3$ dimensions, $N = 8{,}000$ random locations on the sphere ($\mathbf{S}_3^3$), and the Matérn covariance function with $\rho = 10$, $\nu = 0.5$ (exponential), Total Degree index Set $\Lambda(w)$ with $w = 4$, and $p = 35$. As observed from (a) and (b) most of entries of the matrix $\mathbf{C}_W(\boldsymbol{\theta})$ are very small.

$t, t-1, \ldots$ are shown. The polynomial space of the index set $\Lambda(w)$ is restricted to the Hyperbolic Cross on a n-Sphere with $d = 50$ dimensions. The domain decomposition is formed with a RP tree. The level of the index set is set to $w = 5$, which corresponds $p = 1426$. The covariance function is Matérn with $\nu = 1.25$, $\rho = 10$. The distance criterion is set to

$$\tau_{i,j} := 2^{(t-i)/2} 2^{(t-j)/2} \tau,$$

for $i = 1, \ldots, t$ and $j = 1 \ldots, t$, where $\tau = 1 \times 10^{-6}$ or $\tau = 1 \times 10^{-7}$.

The first observation to notice is that all the sparse matrices $\tilde{\mathbf{C}}_W^i(\boldsymbol{\theta})$, $i = t, t-1, \ldots$ *are very well conditioned, thus numerically stable.* This is in contrast to the original covariance matrices that are poorly conditioned, as are shown in Table 6. The sparsity of $\tilde{\mathbf{C}}_W^i(\boldsymbol{\theta})$ and the Cholesky factor $\mathbf{G}$ are shown in columns 8 and 10. The construction time t_{con} is shown in column 9.

It is noted that the sparse matrices in Table 3 are built with a direct summation method. Although it appears that the multi-level covariance matrices are constructed close to linear time, the computational cost should approach quadratic time for a large enough number of observation nodes and $0 \leq i << t$. This computational cost will be further improved as better high dimensional fast summation techniques are developed.

Table 3: Sparsity test on the sparse matrices $\tilde{\mathbf{C}}_W^i$, $i = t, t-1, \ldots$. The polynomial space of the index set $\Lambda(w)$ is restricted to the Hyperbolic Cross on a n-Sphere with $d = 50$ dimensions. The domain decomposition is formed from a RP tree. The level of the index set is $w = 5$, which corresponds $p = 1426$. The kernel function is Matérn with $\nu = 1.25$, $\rho = 10$. The first column is the number of random n-Sphere nodes. The second is the maximum level of the RP tree and i is the level of the sparse matrix $\tilde{\mathbf{C}}_W^i$. The forth column is the size of the matrix $\tilde{\mathbf{C}}_W^i$ and τ is the distance criterion parameter. The sixth column, t_{ML}, is the total time for the construction of the multi-level basis. The seventh column is the sparsity of $\tilde{\mathbf{C}}_W^i$. The eighth column, t_{con} is the total time for the construction of the matrix $\tilde{\mathbf{C}}_W^i$. The ninth column is the sparsity of the Cholesky factor $\mathbf{G}$ (with nested dissection reordering) of the sparse matrix $\tilde{\mathbf{C}}_W^i$. The last column is the total time to compute the Cholesky factor $\mathbf{G}$.

N	t	i	$\kappa(\tilde{\mathbf{C}}_W^i)$	Size	τ	t_{ML}	$nz(\tilde{\mathbf{C}}_W^i(\boldsymbol{\theta}))$	t_{con}	$nz(\mathbf{G})$	t_{sol}
32,000	3	3	1.88	20,592	10^{-6}	116.8	6.3%	104.4	6.3%	1.8
32,000	3	2	2.84	26,296	10^{-6}	119.2	10.0%	346.0	10.0%	6.2
32,000	3	1	3.85	29,148	10^{-6}	118.8	13.5%	876.2	13.8%	12.5
32,000	3	0	5.02	30,574	10^{-6}	118.8	16.9%	1386.3	17.1%	19.3
64,000	4	4	1.93	41,184	10^{-7}	290.7	3.1%	211.3	3.1%	3.6
64,000	4	3	2.39	52,592	10^{-7}	310.8	4.3%	533.4	4.3%	10.0
64,000	4	2	3.21	58,296	10^{-7}	287.8	5.9%	1143.8	5.9%	20.3
64,000	4	1	5.08	61,148	10^{-7}	285.8	7.6%	2697.0	7.6%	87.6
128,000	5	5	2.01	82,368	10^{-7}	632.9	1.6 %	614.7	1.6 %	17.2
128,000	5	4	2.46	105,184	10^{-7}	660.5	2.2 %	1390.1	2.2 %	61.6

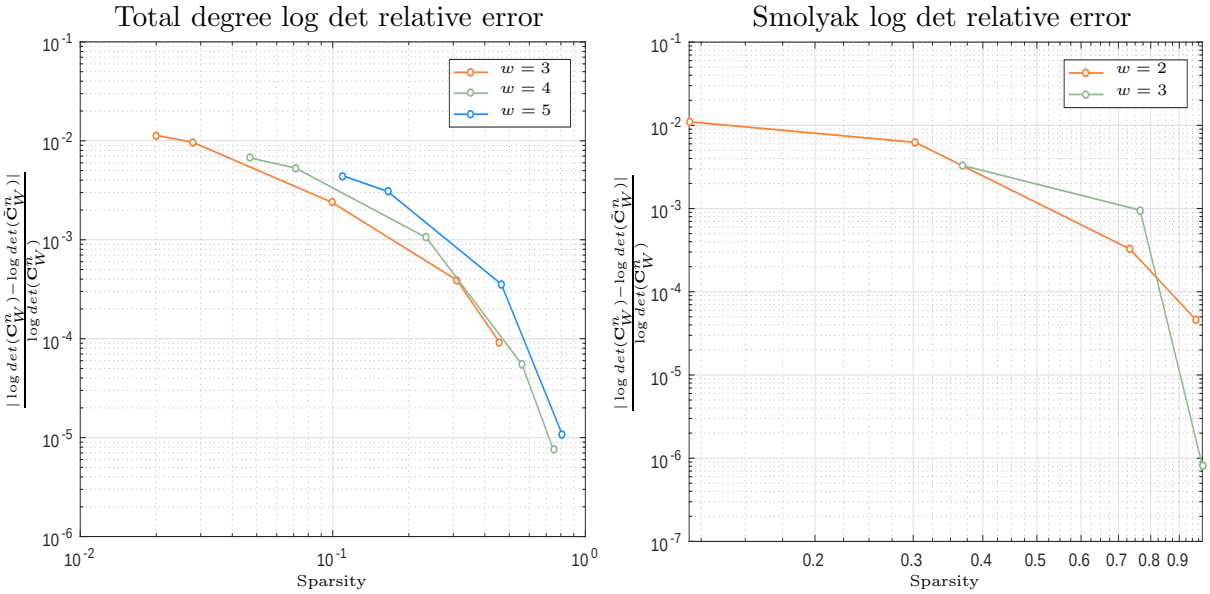

(a) TD, $N = 16{,}000$, $d = 5$, kD-tree (b) SM, $N = 16{,}000$, $d = 10$, kD-tree

Figure 9: Relative log determinant error $\frac{|\log\det\tilde{\mathbf{C}}_W(\boldsymbol{\theta}) - \log\det\mathbf{C}_W(\boldsymbol{\theta})|}{|\log\det\mathbf{C}_W(\boldsymbol{\theta})|}$ with respect to the sparsity of $\tilde{\mathbf{C}}_W^n$ from Random hyper-sphere data $\mathbf{S}_3^d$ with N =16,000 observations, Matérn covariance parameters $\nu = 0.5$, $\rho = 10$ and binary kD-tree. The index sets $\Lambda(\omega)$ are chosen from (a) Total Degree and (b) Smolyak index sets.

8.2 Estimation

In this section estimation results are presented for the Matérn covariance matrix on high dimensional n-Sphere random locations by solving multi-level log-likelihood

$$\hat{\boldsymbol{\theta}}_{,k} := \underset{\boldsymbol{\theta}\in(0,\infty)\times(0,\infty)}{\operatorname{argmax}} \quad \hat{\ell}_W(\mathbf{Z}_{W,k,i}^d, \boldsymbol{\theta}),$$

where $\mathbf{Z}_{W,k,i}^d := [\mathbf{W}_t^{\mathrm{T}}, \ldots, \mathbf{W}_i^{\mathrm{T}}]^{\mathrm{T}}\mathbf{Z}_k^d$, for $i = t, t-1, \ldots$.

Observation data $\mathbf{Z}_4^d$ is built from the n-Sphere $\mathbf{S}_4^d$ for $d = 10$ and $d = 50$ dimensions. Similarly, the data set $\mathbf{Z}_5^d$ is generated for $n = 16,000$ observations. The covariance function is Matérn with $\nu = 5/4$ and $\rho = 1$. To test the performance of the multi-level estimator, $M = 100$ realizations of $\mathbf{Z}_4^d$ and $\mathbf{Z}_5^d$ are generated. Note that 32,000 was the maximum generated data sets due to memory limitations.

The optimization problem of the log-likelihood function (14) is solved using a fmincon iteration search for ν and ρ from the optimization toolbox in MATLAB [23]. The tolerance level is set to 10^{-6}.

In Table 4 the mean and standard deviation of the Matérn covariance parameter estimates $\hat{\nu}$ and $\hat{\rho}$ are presented. The mean estimate $\mathbb{E}_M[\hat{\nu}]$ refers to the mean of M estimates $\hat{\nu}$ for the M realizations of the stochastic model. Similarly, $std_M[\hat{\nu}]$ refers to the standard deviation of the M realizations. As i is reduced from t there is a tendency of a drop in the standard deviation $std_M[\hat{\nu}]$ of the estimator $\hat{\nu}$. However, for $i < t-1$ the standard deviation $std_M[\hat{\nu}]$ does not improve significantly.

Table 4: Estimation of parameters $\hat{\nu}$ and $\hat{\rho}$ with i) Total Degree polynomial index set $\Lambda(w)$, ii) (a) kD or (b) RP tree, iii) n-Sphere with (a) $d = 10$ or (b) $d = 50$ dimensions, iv) The parameters of the observation data $\mathbf{Z}$ are formed with the following Matérn parameters: $\nu = 1.25$ and $\rho = 1$, v) The number of realizations of the Gaussian random field model is set to $M = 100$. and vi) The distance criterion τ is set to 10^{-7}. The first to forth columns are self-explanatory. The fifth column is the mean error of $\hat{\nu}$ with M realization. The sixth column is the mean error of $\hat{\rho}$. The last two columns are the standard deviation of M realizations of the parameters $\hat{\nu}$ and $\hat{\rho}$.

(a) Total Degree, kD tree, n-Sphere, $d = 10$, $M = 100$, $\nu = 1.25$, $\rho = 1$, $\tau = 10^{-7}$

N	w	t	i	$\mathbb{E}_M[\hat{\nu} - \nu]$	$\mathbb{E}_M[\hat{\rho} - \rho]$	$std_M[\hat{\nu}]$	$std_M[\hat{\rho}]$
16000	2	7	7	-1.2e-02	1.3e-02	1.1e-01	5.8e-02
16000	2	7	6	-5.2e-02	2.8e-02	6.5e-02	3.8e-02
16000	2	7	5	-1.6e-01	8.8e-02	4.7e-02	3.8e-02
16000	3	5	5	2.8e-02	-1.0e-02	9.8e-02	5.3e-02
16000	3	5	4	-5.2e-02	3.1e-02	6.1e-02	3.9e-02
16000	3	5	3	-1.5e-01	9.0e-02	4.1e-02	3.4e-02
32000	2	8	8	8.5e-03	-1.1e-03	6.7e-02	3.9e-02
32000	2	8	7	-5.0e-02	2.8e-02	4.5e-02	2.9e-02
32000	2	8	6	-1.5e-01	8.4e-02	3.4e-02	2.9e-02
32000	3	6	6	6.6e-04	2.7e-03	6.9e-02	4.2e-02
32000	3	6	5	-6.1e-02	3.8e-02	4.1e-02	2.9e-02
32000	3	6	4	-1.5e-01	9.7e-02	3.0e-02	2.7e-02

(b) Hyperbolic Cross, RP tree, n-Sphere, $d = 50$, $M = 100$, $\nu = 1.25$, $\rho = 1$, $\tau = 10^{-7}$

N	w	t	i	$\mathbb{E}_M[\hat{\nu} - \nu]$	$\mathbb{E}_M[\hat{\rho} - \rho]$	$std_M[\hat{\nu}]$	$std_M[\hat{\rho}]$
16000	2	7	7	-4.1e-02	3.3e-02	2.9e-01	1.0e-01
16000	2	7	6	-5.7e-02	2.5e-02	2.5e-01	6.8e-02
16000	2	7	5	-9.2e-02	2.7e-02	1.9e-01	5.3e-02
16000	3	6	6	-5.4e-03	1.3e-02	2.4e-01	6.5e-02
16000	3	6	5	-7.1e-02	2.3e-02	2.0e-01	5.5e-02
16000	3	6	4	-1.6e-01	4.1e-02	1.6e-01	5.0e-02
16000	4	2	2	-2.3e-02	2.4e-02	2.4e-01	8.8e-02
16000	4	2	1	-3.9e-02	2.0e-02	1.8e-01	6.1e-02
16000	4	2	0	-9.6e-02	3.4e-02	1.4e-01	4.7e-02
32000	2	8	8	1.3e-04	1.1e-02	2.4e-01	6.0e-02
32000	2	8	7	-2.2e-02	1.1e-02	2.0e-01	4.7e-02
32000	2	8	6	-8.0e-02	2.2e-02	1.7e-01	4.5e-02
32000	3	7	7	-2.5e-02	1.4e-02	2.1e-01	5.0e-02
32000	3	7	6	-7.0e-02	2.0e-02	1.6e-01	4.1e-02
32000	3	7	5	-1.7e-01	4.1e-02	1.2e-01	3.9e-02

8.3 Prediction

The multi-level kriging prediction approach is now tested on the numerical solution of the system of equations $\bar{\mathbf{C}}_W(\boldsymbol{\theta})\bar{\boldsymbol{\gamma}}_W = \bar{\mathbf{Z}}_W$ for $d = 3$ and $d = 50$ dimensions.

Pre-conditioned results for computing $\bar{\mathbf{C}}_W(\boldsymbol{\theta})\bar{\boldsymbol{\gamma}}_W = \bar{\mathbf{Z}}_W$ for the hypercube data set with $d = 3$ dimensions, kD tree, and the Total Degree index set $\Lambda(w)$ are shown in Table 5. The Matérn covariance coefficients $\boldsymbol{\theta} = (\nu, \rho)$ are set to (3/4,1/6). The relative error of the residual of PCG method for the unpreconditioned system is set to $\varepsilon = 10^{-3}$. The KIFMM is set to high accuracy.

The construction of the preconditioner $\mathbf{P}_W$ grows initially with quadratic time. However, around 128,000 observations the KIFMM starts to kick in and the rate is reduced to almost linear time, as expected. The condition number of the covariance matrices are fairly large, making this problem somewhat hard to solve numerically. The results show that 512,000 size problems with good accuracy are feasible with a single 4-core processor.

Table 5: Pre-conditioned results for computing $\bar{\mathbf{C}}_W(\boldsymbol{\theta})\bar{\boldsymbol{\gamma}}_W = \bar{\mathbf{Z}}_W$ for the hypercube data set with $d = 3$ and the Total Degree index set $\Lambda(w)$. The Matérn covariance coefficients $\boldsymbol{\theta} = (\nu, \rho)$ are set to (3/4,1/6). The relative error of the residual of PCG method for the unpreconditioned system is set to $\varepsilon = 10^{-3}$. The KIFMM is set to high accuracy. The second column is the condition number of the covariance matrix $\mathbf{C}$. The third column is the number of iterations needed to obtain 10^{-3} relative error of the unpreconditioned system with $\mathbf{C}_W(\boldsymbol{\theta})$. Let itr($\mathbf{C}_W$) be the number of CG iterations needed for convergence for $\epsilon = 10^{-3}$ residual accuracy. The fourth column presents the wall clock times for the preconditioner computation. The PCG iteration wall clock times for $\mathbf{C}_W$ are given in the fifth column. The last column represents the total wall clock time to compute $\boldsymbol{\gamma}_W = \bar{\mathbf{C}}_W(\boldsymbol{\theta})^{-1}\bar{\mathbf{Z}}_W$.

$\boldsymbol{\theta} = (3/4, 1/6)$, $d = 3$, $w = 7$ $(p = 120)$

N	$\kappa(\mathbf{C})$	itr($\mathbf{C}_W$)	$\mathbf{P}_W$ (s)	Itr (s)	Total (s)
16,000	3.4×10^6	91	50	145	195
32,000	3.0×10^7	168	192	459	651
64,000	-	334	784	2,060	2844
128,000	-	334	3,368	5,427	8616
256,000	-	430	9,828	10,043	19871
512,000	-	970	21,618	51,248	72866

In Table 6 the multilevel prediction method is tested on a Gaussian $\exp(-\frac{r^2}{2h^2})$ and Matérn kernels for $d = 50$ dimensions. The first test is for the Gaussian covariance with $h = \sqrt{10}$ using the ASKIT for the matrix vector products. This case gives rise to badly condition covariance matrices that are difficult to solve with an iterative method as shown from the very large condition numbers of the second column. However, the multi-level prediction method is feasible to solve large problems in high dimensions. Note the wall clock time for convergence grows quadratically. This is mostly due to the quadratic performance of ASKIT as the number of observations increases above 32,000. However, it is still better than using a direct approach.

The second test is for the Matérn covariance function with $\nu = 5/4$ and $\rho = 1$. The direct method is used for the matrix vector products as the performance of ASKIT is only just slightly better. The covariance matrices $\mathbf{C}(\boldsymbol{\theta})$ still have large condition numbers. However, convergence of the multi-level covariance matrix requires only a few iterations due to the excellent condition numbers. For this case it is not necessary to compute the pre-conditioner.

The last test is on the Matérn covariance function with $\nu = 5/4$ and $\rho = 5$, which gives a fairly large overlap between the observations locations. The multi-level covariance matrices are still well conditioned and convergence rates only require a few iterations to reach 10^{-3} relative error.

Table 6: Pre-conditioned results for computing $\bar{\mathbf{C}}_W(\boldsymbol{\theta})\bar{\gamma}_W = \bar{\mathbf{Z}}_W$ for the n-sphere data set with $d = 50$ dimensions, Hyperbolic Cross index set, Gaussian and Matérn covariance function. The polynomial index $\Lambda(w)$ corresponds to the Hyperbolic Cross. The accuracy of the unpreconditioned system is set to $\varepsilon = 10^{-3}$. The second column is the condition number of the covariance matrix with respect to the n-Sphere observations $\mathbf{S}_4^{50}$ and $\mathbf{S}_5^{50}$. The prediction ML method is tested on three cases. The forth column is the computational time needed to compute the Multi-level Basis. (a) The first test is on the Gaussian covariance $\exp(-\frac{r^2}{2h^2})$ using the ASKIT Gaussian kernel for the matrix vector products. This case is badly conditioned and difficult to solve, as shown from the condition numbers of the second column. (b) The second test is with the Matérn covariance function with $\nu = 5/4$ and $\rho = 1$. The direct method is used for the matrix vector products. (c) The last test is on the Matérn kernel with $\nu = 5/4$ and $\rho = 5$.

(a) $h = \sqrt{10}$, $d = 50$, HC, $w = 4$ ($p = 1376$), ASKIT

N	$\kappa(\mathbf{C})$	itr($\mathbf{C}_W$)	MB (s)	$\mathbf{P}_W$ (s)	Itr (s)	Total (s)
16,000	1.8×10^9	24	45	105	235	385
32,000	1.2×10^{10}	62	108	338	1775	2221
64,000	-	79	263	1176	9516	10955
128,000	-	134	620	4776	77220	82616

(b) $\boldsymbol{\theta} = (\nu, \rho) = (5/4, 1)$, $d = 50$, HC, $w = 4$ ($p = 1376$), No preconditioner, Direct

N	$\kappa(\mathbf{C})$	itr($\mathbf{C}_W$)	MB (s)	$\mathbf{P}_W$ (s)	Itr (s)	Total (s)
16,000	2.3×10^4	6	45	-	350	399
32,000	4.9×10^4	7	116	-	1660	1276
64,000	-	9	279	-	8370	8649
128,000	-	11	645	-	41006	41651

(c) $\boldsymbol{\theta} = (\nu, \rho) = (5/4, 5)$, $d = 50$, HC, $w = 4$ ($p = 1376$), No preconditioner, Direct

N	$\kappa(\mathbf{C})$	itr($\mathbf{C}_W$)	MB (s)	$\mathbf{P}_W$ (s)	Itr (s)	Total (s)
16,000	2.8×10^6	7	46	-	406	452
32,000	6.4×10^6	9	115	-	2083	2198
64,000	-	11	272	-	10219	10491
128,000	-	14	643	-	51870	52513

9 Conclusions

In this paper a multi-level kriging method is developed that scales well with high dimensions. A multi-level basis is constructed from a RP or kD tree and for the choice of polynomial Total Degree, Smolyak and Hyperbolic Cross index sets $\Lambda(w)$. The approach described in the paper has the following characteristics and advantages:

i) The multilevel method is numerically stable. Hard estimation and prediction of large dimensional problems are now feasible.

ii) The method is efficiently implemented by using a combination of MATLAB, c++ software packages and dynamic libraries.

iii) Numerical results of up to 50 dimensional problems for estimation and prediction. These problems are difficult to solve directly due to the large condition numbers, but feasible with the multi-level method.

iv) A-posteriori error estimates for the solution of $\mathbf{C}_W(\boldsymbol{\theta})\boldsymbol{\gamma}_W = \mathbf{Z}_W$ are derived for extended Smolyak index sets.

v) The efficiency of this approach will be further improved as high dimensional fast summation methods are developed.

vi) A-posteriori estimates for Hyperbolic Cross index sets will be studied in a future paper. This requires deriving $L^\infty(\Gamma^d)$ error estimates for the HC sparse grid.

Acknowledgements: I appreciate the help and advice from Bill March, George Biros and Lexing Ying, for setting up the KIFMM and ASKIT software packages, and from Lorenzo Tamellini for his help and advice on sparse grids. I also appreciate the support that King Abdullah University of Science and Technology has provided to this project.

References

[1] Intel Math Kernel Library. http://software.intel.com/en-us/articles/intel-mkl/.

[2] M. Abramowitz and I.A. Stegun. *Handbook of Mathematical Functions: With Formulas, Graphs, and Mathematical Tables*. Applied mathematics series. Dover Publications, 1964.

[3] I. Babuska, F. Nobile, and R. Tempone. A stochastic collocation method for elliptic partial differential equations with random input data. *SIAM Review*, 52(2):317–355, 2010.

[4] J. Bäck, F. Nobile, L. Tamellini, and R. Tempone. Stochastic spectral galerkin and collocation methods for PDEs with random coefficients: A numerical comparison. In Jan S. Hesthaven and Einar M. Rnquist, editors, *Spectral and High Order Methods for Partial Differential Equations*, volume 76 of *Lecture Notes in Computational Science and Engineering*, pages 43–62. Springer Berlin Heidelberg, 2011.

[5] J. Baglama and L. Reichel. Augmented implicitly restarted lanczos bidiagonalization methods. *SIAM Journal on Scientific Computing*, 27(1):19–42, 2005.

[6] V. Barthelmann, E. Novak, and K. Ritter. High dimensional polynomial interpolation on sparse grids. *Advances in Computational Mathematics*, 12:273–288, 2000.

[7] J. Beck, F. Nobile, L. Tamellini, and R. Tempone. Convergence of quasi-optimal stochastic galerkin methods for a class of {PDES} with random coefficients. *Computers & Mathematics with Applications*, 67(4):732 – 751, 2014. High-order Finite Element Approximation for Partial Differential Equations.

[8] J. E. Castrillón-Candás, M. G. Genton, and R. Yokota. Multi-level restricted maximum likelihood covariance estimation and kriging for large non-gridded spatial datasets. *Spatial Statistics*, 18, Part A:105 – 124, 2016. Spatial Statistics Avignon: Emerging Patterns.

[9] J. E. Castrillón-Candás, J. Li, and V. Eijkhout. A discrete adapted hierarchical basis solver for radial basis function interpolation. *BIT Numerical Mathematics*, 53(1):57–86, 2013.

[10] J. E. Castrillon-Candas, F. Nobile, and R. Tempone. Analytic regularity and collocation approximation for pdes with random domain deformations. *Computers and Mathematics with applications*, 71(6):1173–1197, 2016.

[11] Y. Chen, T. A. Davis, W. W. Hager, and S. Rajamanickam. Algorithm 887: Cholmod, supernodal sparse cholesky factorization and update/downdate. *ACM Trans. Math. Softw.*, 35(3):22:1–22:14, October 2008.

[12] S. Dasgupta and Y. Freund. Random projection trees and low dimensional manifolds. In *Proceedings of the Fortieth Annual ACM Symposium on Theory of Computing*, STOC '08, pages 537–546, New York, NY, USA, 2008. ACM.

[13] T. Davis and W. Hager. Modifying a sparse cholesky factorization. *SIAM Journal on Matrix Analysis and Applications*, 20(3):606–627, 1999.

[14] T. Davis and W. Hager. Multiple-rank modifications of a sparse cholesky factorization. *SIAM Journal on Matrix Analysis and Applications*, 22(4):997–1013, 2001.

[15] T. Davis and W. Hager. Row modifications of a sparse cholesky factorization. *SIAM Journal on Matrix Analysis and Applications*, 26(3):621–639, 2005.

[16] T. A. Davis and W. W. Hager. Dynamic supernodes in sparse cholesky update/downdate and triangular solves. *ACM Trans. Math. Softw.*, 35(4):27:1–27:23, February 2009.

[17] T. Gerstner and M. Griebel. Dimension-adaptive tensor-product quadrature. *Computing*, 71(1):65–87, September 2003.

[18] F. Jabbari. *Chapter 3: Eigenvalues, singular values and pseudo inverse*. University of California Irvine, Irvine, California. http://gram.eng.uci.edu/ fjabbari/me270b/me270b.html.

[19] B. N. Khoromskij, A. Litvinenko, and H. G Matthies. Application of hierarchical matrices for computing the karhunen–loève expansion. *Computing*, 84(1-2):49–67, 2009.

[20] R. M. Larsen. Lanczos bidiagonalization with partial reorthogonalization, 1998.

[21] A. Litvinenko and H. G. Matthies. Sparse data representation of random fields. *PAMM*, 9(1):587–588, 2009.

[22] W. B. March, B. Xiao, and G. Biros. ASKIT: approximate skeletonization kernel-independent treecode in high dimensions. *CoRR*, abs/1410.0260, 2014.

[23] MATLAB. *version 9.0 (R2016a)*. The MathWorks Inc., Natick, Massachusetts, 2016.

[24] V. Minden, A. Damle, K.L. Ho, and L. Ying. Fast spatial gaussian process maximum likelihood estimation via skeletonization factorizations. *ArXiv*, (arXiv:1603.08057v3), 2016.

[25] H. B. Nielsen, S. N. Lophaven, and J. Søndergaard. *DACE - A Matlab Kriging Toolbox*. Informatics and Mathematical Modelling, Technical University of Denmark, DTU, Richard Petersens Plads, Building 321, DK-2800 Kgs. Lyngby, 2002.

[26] F. Nobile, L. Tamellini, and R. Tempone. Convergence of quasi-optimal sparse-grid approximation of hilbert-space-valued functions: application to random elliptic pdes. *Numerische Mathematik*, 134(2):343–388, 2016.

[27] F. Nobile, R. Tempone, and C. Webster. An anisotropic sparse grid stochastic collocation method for partial differential equations with random input data. *SIAM Journal on Numerical Analysis*, 46(5):2411–2442, 2008.

[28] F. Nobile, R. Tempone, and C. Webster. A sparse grid stochastic collocation method for partial differential equations with random input data. *SIAM Journal on Numerical Analysis*, 46(5):2309–2345, 2008.

[29] W. Nowak and A. Litvinenko. Kriging and spatial design accelerated by orders of magnitude: Combining low-rank covariance approximations with fft-techniques. *Mathematical Geosciences*, 45(4):411–435, 2013.

[30] D. Pfluger, B. Peherstorfer, and H. J. Bungartz. Spatially adaptive sparse grids for high-dimensional data-driven problems. *Journal of Complexity*, 26(5):508 – 522, 2010. SI: {HDA} 2009.

[31] S. Smolyak. Quadrature and interpolation formulas for tensor products of certain classes of functions. *Soviet Mathematics, Doklady*, 4:240–243, 1963.

[32] L. Ying, G. Biros, and D. Zorin. A kernel-independent adaptive fast multipole method in two and three dimensions. *Journal of Computational Physics*, 196(2):591–626, 2004.